

Induced Automation: Evidence from Firm-level Patent Data

Antoine Dechezleprêtre David Hémous Morten Olsen
Carlo Zanella*

April 2021

Abstract

Do higher wages lead to more automation innovation? To answer this question, we first use the frequency of certain keywords in patent text to create a new measure of automation innovation in machinery. We show that our measure is correlated with a reduction in routine tasks in a cross-sectoral analysis in the US. We combine macroeconomic data from 41 countries and information on geographical patent history to build firm-specific measures of low- and high-skill wages. In a firm-level panel analysis, we find that an increase in low-skill wages leads to more automation innovation with an elasticity between 2 and 5. Placebo regressions show that the effect is specific to automation innovations. Finally, we focus on a specific labor market shock, the German Hartz reforms, and show that they reduced automation innovations by those non-German firms relatively more exposed to Germany.

JEL: O31, O33, J20

KEYWORDS: Automation, Innovation, Patents, Income Inequality

*Antoine Dechezleprêtre, OECD, David Hémous, University of Zurich, Morten Olsen, University of Copenhagen and Carlo Zanella, University of Zurich. David Hémous gratefully acknowledges the financial support of the European Commission under the ERC Starting Grant 805007 AUTOMATION. We thank Daron Acemoglu, Lorenzo Casaburi, Patrick Gaule, Michael McMahon, Pascual Restrepo, Joachim Voth and Fabrizio Zilibotti among others for helpful comments and suggestions. We also thank seminar and conference participants at the University of Zurich, Swiss Macro workshop, the University of Copenhagen, the TRISTAN workshop in Bayreuth, the University of Bath, London Business School, the NBER Macroeconomics Across Time and Space Conference, the NBER Summer Institute, the AlpMacro Conference, LMU, Oxford University, Helsinki Graduate School of Economics, TSE, Ecares, Collège de France, Berkeley, EIEF, EC-OECD and IMT Lucca. We thank Amedeo Andriollo, Selina Schön and Shi Suo for fantastic research assistance.

1 Introduction

Do higher wages lead to more labor-saving innovations? And if so, by how much? At a time of fast progress in automation technologies and of political campaigns pushing for higher minimum wages, answering these questions is of central importance as the endogeneity of automation innovations affects both the cost of such policy intervention and their long-term effects. Our paper is the first to establish a causal effect of an increase in labor costs on automation innovations at the firm level.

Answering these question requires overcoming two challenges: identifying automation innovations and finding a source of exogenous variation in labor costs from the perspective of innovating firms. To overcome the first challenge, we develop a new classification of automation patents. We think of automation innovations as innovations which allow for the replacement of workers with machines in certain tasks. We focus on patents in machinery to which our identification strategy is ideally suited. Our classification follows a two-step procedure where we first classify technology categories (IPC and CPC codes) using patent text and then patents using their technology categories. This new classification presents a certain number of advantages: it is transparent, it covers a wide range of automation technologies and it can be built at a highly disaggregated sectoral level. Furthermore, we reproduce the cross-sectoral analysis of Autor, Levy and Murnane (2003) but adding our measure of automation. We find that in the United States, sectors which use equipment with a high share of automation saw a large decrease in routine tasks.

At the country level, technology and wages are co-determined. To find exogenous variation in labor costs, we exploit the fact that automation innovators are often equipment manufacturers which sell their machines to downstream firms in various countries. We conduct two separate exercises. First, we expand on the methodology of Aghion, Dechezleprêtre, Hémous, Martin and Van Reenen (2016, henceforth ADHMOV) and use variation in country-level wages. We rely on the PATSTAT database, which contains close to the universe of patents. For each firm, we compute the geographical distribution of its machinery patents pre-sample, which we use as a measure for the distribution of the firm's international exposure. We then compute firm-specific weighted averages of low- and high-skill labor costs using country-level data. These firm-specific labor costs (referred to as wages for simplicity) proxy for the average labor cost paid by the downstream firms of the innovating firms. As a result, for, say, two German firms, we identify the effect of an increase in, say, US wages, on automation innovations, by comparing how

much more automation innovations increase for the firm which has the higher market exposure to the US.

We conduct our main analysis over the sample period 1997-2011 and use wage data for 41 countries with automation patents for 3,341 firms. We find a substantial effect of wages on automation innovations: higher low-skill wages lead to more automation innovations with an elasticity between 2 and 5 depending on specification. Higher high-skill wages tend to reduce automation, a finding in line with the capital-skill complementarity hypothesis (Krusell, Ohanian, Rios-Rull and Violante, 2000). Our results are robust to the inclusion of country-year fixed effects for the innovator home country and the exclusion of the home country from the wage variable. Importantly, our results are specific to automation innovations and do not extend to other innovations in machinery.

In a second exercise, we focus on a specific labor market shock, namely the Hartz reforms in Germany in 2002-2004. The Hartz reforms are credited with increasing labor supply and reducing labor costs notably for low-skill workers. We analyze the effect of the Hartz reforms and find that they reduced the relative amount of automation innovation undertaken by foreign firms highly exposed to Germany, both in levels and relative to non-automation innovations in machinery.

The theoretical argument that higher wages should lead to more labor-saving technology adoption or innovation is well-understood (e.g. Zeira, 1998). In Hémous and Olsen (forthcoming) and Acemoglu and Restrepo (2018a), wages affect the direction of innovation which can take the form of automation or the creation of new tasks.

There is an extensive empirical literature on the effects of technological change on wages and employment,¹ but the literature on the reverse question is much more limited. A few papers show that labor market conditions affect labor-saving technology adoption in agriculture (Hornbeck and Naidu, 2014, and Clemens, Lewis and Postel, 2018), or manufacturing (Lewis, 2011). Lordan and Neumark (2018) find that minimum wage hikes displace workers in automatable jobs. Unlike these papers our focus is on innovation instead of adoption. This matters because the economic drivers of innovation may differ from those of adoption: innovation may respond differently to macroeconomic variables such as wages; and knowledge spillovers are likely to play a greater role.

Regarding innovation, Acemoglu and Restrepo (2018b) find a positive correlation in

¹See e.g. Autor et. al. (2003), Autor and Dorn (2013) or Gaggl and Wright (2017) for IT, Doms, Dunne and Totske (1997) for factory automation, Graetz and Michaels (2017) or Acemoglu and Restrepo (2020) for robots, Mann and Püttmann (2018) or Bessen, Goos, Salomons and van den Berge (2019) for broader measures of automation.

cross-country regressions between aging and patenting in robotics and numerical control (though their main focus is on adoption). Our paper differs in three ways: first, we build a broader measure of automation innovation in machinery; second, we are interested in the effect of all wage variations not just those arising from demographic trends; and third and most importantly, we conduct our analysis at the firm level instead of the country-industry. Bena and Simintzi (2019) show that firms with a better access to the Chinese labor market decrease their share of process innovations after the 1999 U.S.-China trade agreement.² Andersson, Karadja and Prawitz (2020) look at the effect of emigration to the US in the 19th century in Sweden and find that more exposed municipalities experienced an increase in innovation (but they do not identify automation innovations). In a paper subsequent to ours, Danzer, Feuerbaum and Gaessler (2020) exploit German immigrant settlement policy to show that increases in labor supply discourage local automation innovation, while we exploit firm-level variation and focus on the effect of labor cost on global innovation.

A large literature shows that the direction of innovation is endogenous in other contexts (e.g. Acemoglu and Linn, 2004, and Popp, 2002). Here, we build on ADHMOV, who use firm-level variations in gas prices to show that higher gas prices lead firms in the auto industry to engage more in clean and less in dirty innovations.³

In contemporaneous work, Mann and Püttmann (2020) use machine-learning techniques to classify automation patents and Webb (2020) uses a dictionary approach similar to ours to identify robot, software and artificial intelligence patents. We compare our approaches below.

Section 2 contains our first contribution: a classification of automation technologies. Section 3 introduces a simple model to motivate the analysis. Section 4 describes the data and our empirical strategy. Section 5 contains the results of the main analysis on the effect of wages on automation innovations. Section 6 discusses the event study of the Hartz reforms. Section 7 concludes. The Appendix provides additional robustness checks and details on our methodology.

²Process innovations and automation innovations are not the same: some process innovations reduce other costs than labor (say, materials costs) and many automation innovations are product innovations (a new industrial robot is a product innovation for its maker).

³Three other papers have used ADHMOV's methodology: Noailly and Smeets (2015) on innovation in electricity generation, Coelli, Moxnes and Ulltveit-Moe (2020) on the effect of trade policy on innovation and Aghion, Bénabou, Martin and Roulet (2020) on the role of environmental preferences and competition in innovation in the auto industry. We methodologically extend this work by including country-year fixed effects and separating the foreign variables.

2 Classifying Automation Patents

In this section, we describe the patent data and our method for classifying automation patents. We then show that our measure of automation predicts a decline in routine tasks (reproducing the analysis of Autor et. al., 2003).

2.1 Our approach to classify patents

Our goal is to identify automation innovations in machinery: that is innovations embedded in equipment goods, such as machine tools or robots, which allow for the replacement of workers in certain tasks. Non-automation innovations, in contrast, may improve energy efficiency, reduce the costs of producing certain machines or increase reliability. We employ a dictionary method on patent data and proceed in three steps: i) We use the existing literature to identify keywords related to automation. ii) For each technology category in machinery (based on the IPC and CPC codes in patent data), we compute the share of patents at the European Patent Office (EPO), which contain one of our automation keywords. We use this measure to classify technology categories as automation or not. iii) We then classify worldwide patents as automation or not depending on whether they belong to an automation technology category.

This strategy of first classifying technology categories and then patents has two advantages over classifying patents directly. First, it allows us to include non-EPO patents in our analysis, for which our main data source (PATSTAT) does not have the text.⁴ More generally, other researchers can now use our technology category classification to classify patents without text and future patents. Second, the IPC and CPC codes (henceforth C/IPC codes),⁵ which we use to define our technology categories, are by themselves informative of the characteristics of patents. The particular wording of a patent is also a signal of these characteristics but patents are written in different styles, and rarely expand on the purpose of the invention, so that the same innovation can often be described with or without using our keywords. Conversely, if a patent uses one of our keywords but does not belong to any C/IPC code where this is common, the inclusion of this keyword is frequently uninformative about the nature of the innovation. That

⁴To give an idea of the increase in sample, over the period 1997-2011 there are 3.19 million patent families with patent applications in at least two offices (a condition we will impose in our main analysis). Among those only around 740,000 have an EPO patent with a description in English.

⁵The IPC is the International Patent Classification and the CPC the Cooperative Patent Classification used by the USPTO and the EPO. The CPC is an extension of the IPC and contains around 250,000 codes in its most disaggregated form.

is, the wording of a given patent is a weak signal of whether that patent corresponds to automation but the *combined* wording of many patents gives a strong signal of whether a technological code corresponds to automation. Therefore, our strategy assumes that technology categories are a better signal of whether a patent is an automation patent or not than the presence of our keywords. Yet, as we do not know which technology categories correspond to automation, we use the text of a subset of patents to classify these first. As a matter of fact, the World Intellectual Property Organization (WIPO) offers on its website a simple tool based on a similar principle: a search engine allows one to identify up to 5 IPC codes most likely to correspond to a set of keywords using the text of the patents in its database.

Alternatively, we could have read and classified a subset of patents and then used machine-learning techniques to classify other patents or technology categories based on patent text. This is the procedure in Mann and Püttmann (2018), whose results we discuss in Section 2.4 and Appendix A.3. Relying on keywords instead of a training set of patents presents several advantages. First, manually classifying patents as automation is a difficult task which cannot be easily systematized and outsourced: often looking at a single patent in isolation is not enough, and one needs to look at several patents within the same technological group to find patterns suggesting that a patent is likely an automation patent. Second, patents are written in a technical language and do not primarily discuss the goal of an innovation, so that only a few words within the text are informative and a machine-learning algorithm would require a very large training set. Third, by using a few keywords instead of a large training set, our approach is more transparent, easily replicable and modifiable and, as researchers, we have fewer degrees of freedom since we pick most of our keywords from the literature.

2.2 Patent data

We use two patent databases maintained by the EPO. For most of our empirical analysis, we use the World Patent Statistical Database (PATSTAT) from Autumn 2018 which contains the bibliographical information of patents from 90 patent-issuing authorities (covering nearly all patents in the world) but not the text of individual patents. Since text analysis is essential to our approach, we supplement with the EP full-text database from 2018, which contains the full text of EPO patent applications.

PATSTAT allows us to identify “patent families”, a set of patent applications across different patent offices which represent the same innovation. For each patent family, we

Table 1: Choice of automation keywords

Keywords	Comments	Source
Automat*	<i>Automation, automatization</i> or <i>automat* at least 5 times</i> or (automat* or autonomous) with (secondary words or warehouse or operator or arm or convey* or handling or inspect* or knitting or manipulat* or regulat* or sensor or storage or store or vehicle system or weaving or welding) in the same sentence at least twice	Own / Doms, Dunne and Troske (DDT) / Acemoglu and Restrepo (AR)
Robot*	Not surgical or medical	DDT and AR
Numerical Control	CNC or numeric* control* or (NC in the same sentence as secondary words)	DDT and AR
Computer-aided design and manufacturing	Computer-aided/-assisted/-supported in the same patent as secondary words CAD or (CAM and not "content addressable memory") in same sentence as secondary words	DDT
Flexible manufacturing		DDT
Programmable logic control	Programmable logic control or [PLC and not (powerline or "power line")]	DDT
3D printer	"3D print*" or "additive manufacturing" or "additive layer manufacturing"	Own
Labor	Including laborious	Own
Secondary words	Machine or manufacturing or equipment or apparatus or machining	

Notes: "In the same sentence as control words" refers to at least one control word. Keywords include i) natural adjacent words (i.e. numerical control includes NC, numerically controlled and numeric control), ii) British/American spelling (i.e. labour/labor) and iii) hyphenated adjectives (i.e. computer aided / computer-aided design). We added words in italics, the others come from AR or DDT. See Appendix for details.

know the date of first application (which we use as the year of an innovation), the patent offices where the patent is applied for, the identity of the applicants and the inventors and the number of citations received by the patent family. To identify the technological characteristics of patents we use their C/IPC codes. Importantly each patent usually has several C/IPC codes. The C/IPC codes form a hierarchical classification systems. Certain types of technologies (for instance fossil fuel engines) can readily be identified to existing groupings of C/IPC codes. Such a grouping does not exist for automation and our goal here is to create one.

2.3 Choosing automation keywords

To tie our hands, we choose most of our keywords from the automation technologies identified in Doms, Dunne and Troske (DDT, 1997) and Acemoglu and Restrepo (AR, 2018b) and complement with a few additional words as described below.⁶ In fact, most of our search terms (for simplicity "keywords") correspond to the co-occurrence of our several words in the same sentence or patent, or the repetition of these words a sufficient number of times. Table 1 describes the list of our search terms together with their origin.

⁶Doms, Dunne and Troske (1997) measure automation using the Survey of Manufacturing Technology (SMT) from 1988 and 1993 conducted by the US Census. The survey asked firms about their use of certain automation and information technologies. Acemoglu and Restrepo (2018b) include imports of automation technology and associate specific HS-categories from Comtrade with automation technology.

We have eight categories of keywords. Five of these, robot*, numerical control, computer-aided design and manufacturing, flexible manufacturing and programmable logic control are automation technologies in DDT or AR. Simply applying these words may result in false positives. For instance “NC” can refer to either “numerical control” or “North Carolina”. To address this issue, we require that these words are either in the same patent or the same sentence as a list of secondary words, such as machinery or equipment, which indicate that the text describes a machine. Furthermore, we add “automation” and “automatization”. The stem “automat*” gather too many false positives such as “automatic transmission”. We resolve this in two ways: either by restricting attention to patents where the frequency is 5 or more or by combining automat* with our secondary words or other words which largely come from technologies described in DDT or AR and often describe tasks (such as manipul*, regulat* or inspect*). We count patents where automat* and one of these words appear in the same sentence at least twice. Finally, we add 3D printing, which was in its infancy when DDT was written, and “labor” which often indicates that an innovation reduces labor costs. The most important keywords are those associated with “automat*” (see Appendix A.2) and Section 5.6 shows that our main results are robust to only using those.

2.4 Automation technology categories and patents

Defining machinery C/IPC codes. We base our classification on the set of EPO patent applications from 1978 till 2018 with a description in English (1,538,370 patent applications), which we denote Ω_{EPO} . We use the keywords to associate technology categories, and not patents directly, to automation. These technology categories are defined as: 6-digit C/IPC codes, all pairs of 4-digit C/IPC codes and, inspired by Aschhoff et al. (2010), pairs combining the union of the 3 digit codes G05 and G06 with any 4-digit C/IPC codes.⁷ The code G05 corresponds to “controlling; regulating” and G06 to “computing; calculating; counting” and they use these combinations to identify advanced manufacturing technologies. We restrict attention to categories which contain at least 100 patents (we group 6-digit codes with the same 4-digit code and less than

⁷Technically, the structure of the C/IPC classification is as follows: C/IPC “classes” have 3 digit codes (for instance B25: “hand tools; portable power-driven tools; handles for hand implements; workshop equipment and manipulators”), “subclasses” have 4 digit codes (for instance B25J: “manipulators; chambers provided with manipulation devices”), main groups have 5 to 7 digit codes (for instance B25J 9: “programme-controlled manipulators”). In the following, we slightly abuse language and refer to classes, subclasses and main groups as 3 digit, 4 digit and 6 digit codes respectively.

100 patents in Ω_{EPO} in common artificial 6 digit codes).

Our keywords are best associated with automation in equipment and we accordingly restrict attention to C/IPC codes which belong to certain technological fields. There are 34 technological fields (see Figure A.1) and we focus on “machine tools”, “handling”, “textile and paper machines” and “other special machines” with some adjustments, which we refer to as “machinery” patents (we use machinery and equipment interchangeably).⁸ This leaves us with 1009 6-digit C/IPC codes. For pairs of 4 digit C/IPC codes or pairings of 4 digit C/IPC codes with G05 or G06 we classify them as belonging to machinery if at least a 4 digit code belongs to that field.

Formally, a patent p is associated with a set of C/IPC codes C_p recorded at a highly disaggregated level. Define the functions $S_3(C_p)$ which extracts the set of unique 3 digit codes in C_p ; $S_4(C_p)$ which extracts the set of unique 4 digit codes and $S_6(C_p)$ which extracts the set of unique 6-digit codes (grouping 6-digit codes with less than 100 patents in Ω_{EPO} at the 4 digit level). Further, define the function $m_4(c_4)$ which takes the value 1 if a 4 digit C/IPC code belongs to the machinery technological fields and 0 otherwise and similarly $m_6(c_6)$ for a 6-digit code. We then define the broad set of machinery technology categories of a patent p , $\widetilde{MT}_p$ as

$$\begin{aligned}\widetilde{MT}_p = & \{c_6 \in S_6(C_p) | m_6(c_6) = 1\} \\ & \cup \{\{c_4, c'_4\} | c_4, c'_4 \in S_4(C_p) \wedge c_4 \neq c'_4 \wedge (m_4(c_4) = 1 \vee m_4(c'_4) = 1)\} \\ & \cup \{\{c_3, c_4\} | \{c_3, c_4\} \in \{S_3(C_p), S_4(C_p)\} \wedge m_4(c_4) = 1 \wedge c_3 \in \{G05, G06\}\}.\end{aligned}$$

We exclude technological categories with less than 100 patents in Ω_{EPO} and denote the remaining set MT_p . The overall set of machinery technology categories is $\mathcal{MT} = \cup_{p \in \Omega_{EPO}} MT_p$.

Defining automation C/IPC codes. A patent is also associated with a text T_p , for each keyword category (automat*, robot, CNC, etc.) we define functions $k^{automat*}(T_p)$, $k^{robot}(T_p)$, $k^{CNC}(T_p)$, etc. which take value 1 if one of the associated keyword is in the text

⁸We exclude F41 and F42 which correspond to weapons and ammunition and are in “other special machines”. Otherwise drones and missiles show up as a highly automated technology. Moreover, we include B42C which corresponds to machines for book production and B07C which corresponds to machines for postal sorting as both correspond to equipment technologies and contain 6-digit codes with a high prevalence of automation keywords; the 6-digit code G05B19 which corresponds to “programme-control systems” and contains a large number of computer numerically controlled machine tool patents without C/IPC from the machine tools technological field; and the 6-digit code B62D65 which deals with engine manufacturing even though the rest of the B62D code deals with the vehicle parts themselves. We verify that these additional codes do not qualitatively affect our results.

Table 2: Examples of 6-digit C/IPC codes in machinery

Code	Description	# patents	Any	Rank	Robot	Automat*	CNC	labor
High prevalence								
B25J5	Manipulators mounted on wheels or on carriages	504	0.91	1	0.87	0.27	0.01	0.1
B25J9	Programme-controlled manipulators.	2809	0.86	4	0.79	0.29	0.07	0.08
B23Q15	Automatic control or regulation of feed movement, cutting velocity or position of tool or work.	591	0.79	7	0.09	0.36	0.65	0.06
A01J7	Accessories for milking machines or devices.	395	0.77	9	0.62	0.52	0	0.1
G05B19	Programme-control systems.	7133	0.7	16	0.22	0.39	0.25	0.08
B65G1	Storing articles, individually or in orderly arrangement, in warehouses or magazines	1064	0.58	29	0.18	0.46	0.01	0.11
Low prevalence								
B23P6	Restoring or reconditioning objects.	613	0.26	266	0.07	0.06	0.05	0.09
A01B63	Lifting or adjusting devices or arrangements for agricultural machines or implements.	264	0.24	306	0.01	0.2	0	0.04
B66D3	Portable or mobile lifting or hauling appliances	215	0.13	677	0.02	0.07	0	0.06

Note: Prevalence of automation keywords for a few 6 digit C/IPC codes. "Any" is the share of patents with any of the keywords. "Rank" is the rank of the code among 1009 6-digit C/IPC codes in machinery with at least 100 patents. "Robot", "Automat*", "CNC" and "labor" are the shares of patents with at least one keyword from these categories.

and 0 otherwise. We define $k^{any}(T_p) = \max \{k^{automat*}(T_p), k^{robot}(T_p), k^{CNC}(T_p), etc.\}$ which takes value 1 if any of the automation keywords are present. For all machinery technology category $t \in \mathcal{MT}$, we define the prevalence of automation keywords $p(t)$ as the share of patents containing at least one of our keywords:

$$p(t) = \frac{\sum_{p \in \Omega_{EPO}} 1_{t \in MT_p} k^{any}(T_p)}{\sum_{p \in \Omega_{EPO}} 1_{t \in MT_p}}.$$

We similarly define the prevalence of specific keyword categories. We extensively checked the C/IPC codes and sampled patents from each category to ensure that the procedure delivered reasonable results and adjusted the keywords accordingly. However, we never modified the classification after carrying out any of our regressions.

Table 2 gives some examples of 6-digit C/IPC codes in machinery with their prevalence of automation keywords $p(t)$ and their rank according that measure. It also shows the prevalence of the most important subcategories (automat*, robots, CNC and labor). C/IPC codes associated with robotics (B25J) have the highest prevalence numbers with up to 91% patents in B25J5 which contain at least one of the keywords. There are also codes associated with machine tools at the top of the distribution such as B23Q15 and codes associated with devices used in the agricultural sector such as A01J7. The last three C/IPC codes are examples with a low prevalence of automation keywords: machine-tools and processes for repairing or reconditioning objects (B23P6), devices typically mounted on tractors (A01B63), and lifting or hauling appliances such as hoists




Figure 1: Prevalence of automation keywords for C/IPC 6 digit codes in machinery

(B66D3). The table also shows that the different sub-measures do not capture the same technologies: the robotic codes are ranked highly thanks to their share of patents with the word “robot”, B23Q15 is high because a lot of patents contain words related to CNC, and B65G1, because a lot of patents contain words associated with automation directly.

Figure 1 gives the histogram of the prevalence of automation keywords for all C/IPC 6 digit codes in machinery. It shows that most C/IPC codes have a low prevalence of automation keywords but a few codes have a very high value. Appendix A.2 gives additional statistics on the prevalence measures.

We define automation technology categories as those with a prevalence measure above some threshold. As our baselines, we choose thresholds at the 90th and 95th percentiles of distribution of the 6 digit code distribution (within machinery), which are given by 0.386 and 0.477, respectively.⁹ Therefore a technology category t belongs to the set of auto90 categories T^{90} if $p(t) > 0.386$ and to the set of auto95 category T^{95} if $p(t) > 0.477$. We then define a patent as an automation patent if it belongs to at least one automation technology category. That is we classify a patent family p from the PATSTAT dataset $\Omega_{PATSTAT}$ as an auto90 patent if $\exists t_p \in MT_p$ such that $t_p \in T^{95}$, and similarly for an auto90 patent. Note that close to 80% of automation patents are identified by the 6 digit alone (see Appendix A.2).¹⁰

⁹Choosing different thresholds is easy and we investigate how robust our results are in Section 5.6.

¹⁰In Appendix A.2.3, we show that the technology categories with a high prevalence of automation keywords remain the same throughout the period considered. In particular, the correlation between the

(19)	 Europäisches Patentamt European Patent Office Office européen des brevets	 (11) EP 2 604 550 B1	Description
(12)	EUROPEAN PATENT SPECIFICATION		OBJECT OF THE INVENTION
(45)	Date of publication and mention of the grant of the patent: 01.10.2014 Bulletin 2014/40	(51) Int Cl.: B65G 1/137 (2006.01) B65G 1/04 (2006.01) B66F 9/07 (2006.01) A47B 96/02 (2006.01) B65G 1/08 (2006.01)	
(21)	Application number: 10855839.6	(86) International application number: PCT/ES2010/070549 (87) International publication number: WO 2012/020149 (16.02.2012 Gazette 2012/07)	
(22)	Date of filing: 12.08.2010		
(54)	AUTOMATIC PLANT FOR STORING AND DISPENSING GOODS AUTOMATISCHE ANLAGE ZUR AUFBEWAHRUNG UND AUSGABE VON WAREN INSTALLATION AUTOMATIQUE POUR STOCKER ET DISTRIBUER DES PRODUITS		
(84)	Designated Contracting States: AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO SE SI SK SM TR	• GONZÁLEZ LÓPEZ, Isabel E-47012 Valladolid (ES)	
(43)	Date of publication of application: 19.06.2013 Bulletin 2013/25	(74) Representative: Ungria López, Javier c/o UNGRIA Patentes y Marcas, S.A., Avda. Ramon y Cajal, 78 28043 Madrid (ES)	
(73)	Proprietor: Automatismos Y Montajes Industriales J. Martin, S.L. 47012 Valladolid (ES)	(56) References cited: EP-A1- 2 113 473 CH-A5- 680 434 DE-A1- 4 336 885 DE-A1- 4 339 055 DE-A1- 19 635 396 DE-A1- 19 724 378 DE-U1- 20 021 440 US-A- 3 782 565 US-A1- 2010 168 910	
(72)	Inventors: • MARTÍN DE PABLO, Francisco Javier E-47012 Valladolid (ES)		

Figure 2: Example of an automation patent

Figure 2 shows an automated storage cabinet patent. We classify it as automation because it contains the 6 digit code B65G 1 which has a high prevalence measure (0.58, see Table 2). This patent itself contains several keywords: a sentence with the words “automatic” and “storing,” and another sentence with “robot”. Appendix Figure A.2 shows an automation patent of a similar storage cabinet that belongs to the same C/IPC code but does not contain any keywords and still describes a labor-saving innovation. The supplemental material on our website provides more examples.

Comparison with Mann and Püttmann (2018). Mann and Püttmann (2018) also classify patents as automation versus non-automation. Our approaches differ in three ways. First, they classify all patents while we focus on machinery. Second, they manually classify a training set and use machine learning to classify US patents in a given period, while we identify technology categories thanks to a dictionary method, so that we (or others) can classify any patent in machinery. Third, they define as automation “a device that carries out a process independently of human intervention”, while we seek to identify innovations that replace workers in existing tasks. Therefore, they classify prevalence measures computed for the first half of the sample and the second half is 0.92.

a number of patents related to elevators and printing machines as automation patents, which we do not (see Appendix A.3 where we compare the two approaches in details).¹¹

2.5 Trends in automation innovations

We generally restrict attention to patent families with applications in at least two countries (referred to as biadic patents). Several studies (e.g. De Rassenfosse et al., 2013, and Dechezleprêtre, Ménière and Mohnen, 2017,) have shown that such patents are of higher quality than others.¹² Focusing on biadic patents is also consistent with our empirical strategy which relies on firms’ exposure to international markets.

Figure 3 plots the evolution of automation biadic patent families. Panel (a) shows that worldwide the share of automation patents in machinery slightly declined between the mid1980s (9.5% in 1985 for auto95) and the mid1990s (7.6% in 1994 for auto95) before increasing quickly (reaching 18.9% in 2015 for auto95 in 2015). Appendix Figure A.3 shows that auto95 patents represent 2.7% of all patents in 2015, a share that has doubled since 1997. It also reports the raw numbers of auto90 and auto95 patents. Figure 3.b shows the trends for auto95 by applicant nationality. The trend for Japan is somewhat distinct: it is initially considerably higher, but declines in the 80s and 90s before picking up in the 2000s though slower than in the other countries. Germany has the highest automation share in 2015.

2.6 Automation and routine tasks

Autor et. al. (2003) (henceforth ALM) show that computerization has been associated with a decrease in routine tasks at the industry level on U.S. data from 1960 to 1998. Here, we briefly analyze how our measure of automation relates to routinization, in part as a way to validate our measure of automation before focusing on our main topic: the effect of wages on induced automation innovation.

¹¹Bessen and Hunt (2007) also use keywords to identify software patents. Webb (2020) focuses on matching three technologies (robotics, software and AI) to the occupations that they may replace. To identify the associated patents, he also uses keywords: he uses the algorithm of Bessen and Hunt (2007) for software patents, while robotics patents are defined as those with “robot” or “manipulat” in the title or abstract but exclude the CPC classes A61 or B01 (to avoid surgical robots). We instead focus on all automation innovation in machinery and since our classification is available at the C/IPC level, it can easily be used and extended by other researchers.

¹²We count applications and not granted patents because in certain patent offices, notably Japan, a patent is only formally granted if the rights of the applicant are challenged. To restrict attention to patent families of even higher quality, we carry out robustness checks where we use patent citations.




Figure 3: Share of automation patents in machinery for biadic families.

As ALM, we run industry level regressions of the type:

$$\Delta T_{jk\tau} = \beta_0 + \beta_C \Delta C_j + \beta_{aut} aut_{j\tau}. \quad (1)$$

$\Delta T_{jk\tau}$ represents the change in tasks of type k in industry j during period τ and ΔC_j is the measure of the change of computerization in sector j (it is computed over the years 1984-1997 and used for all time periods τ). $aut_{j\tau}$ is our patent-based measure of automation intensity in sector j , period τ . We do not first-difference this measure since patenting is already a measure of the flow of knowledge. We take our task measures directly from ALM, and therefore consider 5 types of tasks: nonroutine analytic, non-routine interactive, routine cognitive, routine manual and nonroutine manual. $\Delta T_{jk\tau}$ is measured as 10 times the annual within-industry change in task input measured in percentile of the 1960 task distribution.

To construct $aut_{j\tau}$, we allocate patents in machinery to their sector of use, focusing here on USPTO granted patents. Autor, Dorn, Hanson, Pisano and Shu (2020) match USPTO patents with firm-level data from Compustat and thereby provide detailed sectoral information for corporate patents. We use their data to create a (weighted) concordance table from C/IPC 4 digit codes to 4 digit SIC industries. This mapping can be used to allocate patents to sectors of invention. To get the sector of use, we then combine their mapping with the 1997 capital flow table from the BEA (the capital flow table is similar to an input output table but reports the flows in investment goods in-




Figure 4: Scatter plots of routine tasks changes and automation intensity (auto 95) in 1980-1998 in the United States. The list of sectors is given in Table A.24

stead of intermediate inputs). For each sector j and period τ , we compute $aut_{j\tau}$ as the share of automation patents among machinery patents applied for during this period. We consider 4 time periods: 1970-1980, 1980-1990 and 1990-1998 and the joint time period 1980-1998. We restrict attention to sectors with at least 50 machinery patents over the period. We can then measure automation intensity for 124 sectors in 1980-1998. Our automation measure auto95 is only weakly correlated with computerization with a coefficient of 0.17 (and -0.19 when we weigh industries by employment). See Appendix A.4 for further details on the data.

Figure 4 provides simple scatter plots of the changes in routine tasks and the share of automation patents in machinery over the years 1980-1998 (according to the auto95 definition). There is a clear relationship: sectors with a high share of automation patents experience a larger decline in routine cognitive and routine manual tasks. Given our focus on automation in machinery a decline in routine cognitive tasks might seem surprising at first sight, but several machines replace workers in tasks such as inspection and control (such an example is given in Figure A.2).

Table 3 reports the results of regressions (1) for the auto95 measure. Columns (3) and (4) show that sectors with a high share of automation patents in machinery experienced a large reduction in both cognitive and manual routine tasks in each decade (in the 90s, the coefficient on routine manual tasks is non significant). For instance, Panel B

Table 3: Correlation between changes in task intensity or skill ratio across sectors and automation (auto95)

	(1) Δ Nonroutine analytic	(2) Δ Nonroutine interactive	(3) Δ Routine cognitive	(4) Δ Routine manual	(5) Δ Nonroutine manual
Panel A: 1970 - 80, n=115					
Share of automation patents in machinery	-16.51 (37.17)	119.56** (50.18)	-81.66* (46.02)	-62.89 (43.35)	124.64*** (43.09)
Δ Computer use 1984 - 1997	6.38 (4.98)	10.17 (6.72)	-8.04 (6.16)	-7.33 (5.80)	-4.91 (5.77)
Intercept	0.90 (2.29)	-4.04 (3.10)	5.63** (2.84)	5.64** (2.67)	-6.53** (2.66)
R ²	0.02	0.06	0.04	0.03	0.08
Weighted mean Δ	1.22	4.27	-0.25	0.89	-0.79
Panel B: 1980 - 90, n=115					
Share of automation patents in machinery	88.81*** (29.21)	82.67** (36.90)	-195.57*** (29.78)	-138.69*** (31.89)	35.49 (29.50)
Δ Computer use 1984 - 1997	18.84*** (5.56)	21.39*** (7.02)	-17.70*** (5.67)	-11.54* (6.07)	-1.86 (5.61)
Intercept	-8.39*** (2.71)	-5.64 (3.42)	16.62*** (2.76)	11.53*** (2.96)	-3.74 (2.73)
R ²	0.14	0.10	0.30	0.15	0.02
Weighted mean Δ	2.14	4.89	-2.01	-1.48	-1.33
Panel C: 1990 - 98, n=115					
Share of automation patents in machinery	21.71 (31.82)	42.74 (36.19)	-77.53** (36.65)	-99.72** (38.65)	40.21 (26.37)
Δ Computer use 1984 - 1997	12.97** (5.36)	16.30*** (6.10)	-16.91*** (6.18)	-29.51*** (6.51)	9.40** (4.44)
Intercept	-1.63 (2.83)	-2.31 (3.22)	5.90* (3.26)	9.86*** (3.43)	-5.28** (2.34)
R ²	0.05	0.06	0.08	0.17	0.05
Weighted mean Δ	2.54	4.10	-3.26	-3.43	-0.41

Standard errors are in parentheses. Columns (1) to (5) of Panels A to C each presents a separate OLS regression of ten times the annual change in industry-level task input between the endpoints of the indicated time interval (measured in centiles of the 1960 task distribution) on the share of automation patents in machinery (defined with the 95th percentile threshold) and the annual percentage point change in industry computer use during 1984 - 1997 as well as a constant. In Column (6), the dependent variable is the ratio of high-skill (college graduates) to low-skill (high-school graduates and dropouts) workers. Estimates are weighted by mean industry share of total employment in FTEs over the endpoints of the years used to form the dependent variable. * p<0.1; ** p<0.05; *** p<0.01

indicates that a 1 pp increase in the share of automation patents is associated with a 2 and 1.4 centiles decrease in routine cognitive and manual tasks in the 80s. The standardized beta coefficients are larger than for computerization since they correspond to 2.7 and 1.9 centiles in routine cognitive and manual tasks versus 1.3 and 0.8 though the effect of computerization is larger in the 90s.¹³ We obtain similar results when we restrict attention to biadic patents. In Appendix A.4, we also get similar results but with a lower magnitude when we compute the share of automation patents with the auto90 measure. We also allocate patents to their sector of invention and include this measure as a control in our regressions. We still find a negative effect of the use of automation technologies on routine tasks.

Our classification of machinery patents as automation or non-automation is a prerequisite for our main empirical exercise where we focus on the effect of an increase in labor costs on the innovations of equipment producers. Importantly, given a mapping between C/IPC codes and sectors, it can also deliver a measure of automation at a more detailed sectoral level than alternatives such as robotization. And we have now showed that this measure is uncorrelated with computer use but is associated with a reduction in routine tasks at the sectoral level.

3 A Simple Model

Before carrying out our empirical analysis, we present a simple model to clarify our argument. The model is motivated by the business structure of the largest automation innovators. In 2018, Siemens, the biggest innovator in our sample, had 31% of its workforce in Germany but only 14% of its revenues from there. Its strongest growing division was the Digital Factory Division which provides a broad range of automation technology to manufacturers across the globe. The annual report describes how “The Digital Factory Division offers a comprehensive product portfolio and system solutions for automation technologies used in manufacturing industries, such as automation systems and software for factory automation, industrial controls and numerical control systems, motors, drives and inverters and integrated automation systems for machine tools and production machines...”. Note that this sentence includes a lot of our keywords. The report is centrally interested in how “Changes in customer demand [for automation technology by down-

¹³The employment-weighted standard deviation in the share of automation patents in the 80s for the included industry is 1.4% and the mean 7.8%, while the standard deviation for computerization is 0.07. Meanwhile routine tasks decline by 2 and 1.5 centiles on average for these sectors.

stream manufacturers] are strongly driven by macroeconomic cycles”. Interestingly, it never mentions “cost of labor” as a reason for automation, but instead uses a number of euphemisms such as “increase competitiveness”, “enhance efficiency”, “improve cost position” and “streamline production”. Siemens further discusses how such macroeconomic trends affect its *R&D* decisions.

We incorporate these business features into a model built on Hémous and Olsen (forthcoming). A manufacturing good is produced with a continuum of intermediate inputs according to the Cobb-Douglas production function $Y = \exp\left(\int_0^1 \ln y(i) di\right)$, where $y(i)$ denotes the quantity of intermediate input i . The manufacturing good is the numéraire. Each intermediate input is produced competitively with high-skill labor ($h_{1,i}$ and potentially $h_{2,i}$), low-skill labor, l_i , and potentially machines, x_i , according to:

$$y_i = h_{1,i}^{1-\beta} \left(\gamma(i) l_i + \alpha(i) \nu^\nu (1-\nu)^{1-\nu} x_i^\nu h_{2,i}^{1-\nu} \right)^\beta. \quad (2)$$

$\gamma(i)$ is the productivity of low-skill workers, $\alpha(i)$ is an index which takes the value 0 for non-automated intermediates and 1 for automated intermediates and ν and β are parameters in $(0,1)$. Machines are specific to the intermediate input i . If a machine is invented, it is produced monopolistically, 1 for 1 with the final good so that the monopolist charges a price $p_x(i) \geq 1$. At the beginning of the period, for each non-automated intermediate i , there is an innovator. The innovator creates a machine specific to intermediate i with probability λ if she spends $\theta\lambda^2 Y/2$ units of manufacturing good.

For an automated intermediate input ($\alpha(i) = 1$), the downstream producer is indifferent between using low-skill workers or machines together with high-skill workers in production whenever $w_H^\nu p_x^{1-\nu} = w_L/\gamma(i)$. Therefore, the machine producer is in “Bertrand competition” with low-skill workers. As a machine costs 1, the machine producer charges a price $p_x(i) = \max\{(w_L/\gamma(i))^{\frac{1}{1-\nu}} w_H^{-\frac{\nu}{1-\nu}}, 1\}$, such that machines are used if $w_L/\gamma(i) > w_H^\nu$. Since the manufacturing good is produced according to a Cobb-Douglas production function, we get $p(i)y(i) = Y$ for all intermediates. We can then derive the profits of the machine producer as $\pi_i^A = \max\left(1 - (\gamma(i)/w_L)^{\frac{1}{1-\nu}} w_H^{\frac{\nu}{1-\nu}}, 0\right) \nu\beta Y$.

In turn, at the beginning of the period, the potential innovator solves $\max \lambda \pi_i^A - \theta\lambda^2 Y/2$, giving the equilibrium innovation rate $\lambda = \pi_i^A/(\theta Y)$. As a result, the number of automation innovations is equal to:

$$Aut = \frac{\nu\beta}{\theta} \int_0^1 (1 - \alpha(i)) \max\left(\left(1 - \left(\frac{\gamma(i)}{w_L}\right)^{\frac{1}{1-\nu}} w_H^{\frac{\nu}{1-\nu}}\right), 0\right) di.$$

This expression is increasing in the low-skill wage w_L and decreasing in the high-skill wage w_H . Intuitively, the incentive to replace low-skill workers with machines (and high-skill workers) increases with low-skill wages, leading to a higher demand for machines. The reverse holds for high-skill wages. An upward shift in low-skill worker productivity, $\gamma(i)$, also reduces the number of automation innovations.

To contrast automation with other types of innovations, assume that the production of an intermediate takes place according to:

$$y_i = (q_i m_i)^\delta h_{1,i}^{1-\beta-\delta} (\gamma(i) l_i + \alpha(i) \nu^\nu (1-\nu)^{1-\nu} x_i^\nu h_{2,i}^{1-\nu})^\beta,$$

where m_i denotes non-automation “Hicks” machines with quality q_i . Hicks machines are also produced one-for-one with the final good. Each period one innovator may improve on the available quality of Hicks machines for intermediate i by a factor μ by investing in R&D. If she spends $\theta_m \lambda_m^2 Y/2$ units of the final good, she is successful with probability λ_m . In that case, the innovator becomes the monopolistic provider of Hicks machine i under the pressure of a competitive fringe which has access to the previous technology, and the technology diffuses after one period. Otherwise, the good is produced competitively. The previous analysis on automation innovations remains identical. A successful Hicks innovator can charge a mark-up μ leading to profits $\pi_i^H = (1 - \mu^{-1}) \delta Y$. The innovation rate is then $\lambda_m = (1 - \mu^{-1}) \delta / \theta_m$, so that the number of Hicks innovations is a constant given by $NonAut = \delta (1 - \mu^{-1}) / \theta_m$. In contrast to automation innovations, the number of non-automation innovations is independent of low- or high-skill wages.

4 Empirical Strategy and Data

We now take the predictions of our model to the data. In this section, we present the regression framework and the data construction. Section 5 will discuss results and identification assumptions.

4.1 Empirical strategy

As mentioned above, innovators in automation technologies are often large companies (e.g. Siemens) which sell their automation equipment internationally. Following the logic of our model, the incentives of the downstream producers to adopt automation technology is determined by wages in their local market. As a result, the decision of

innovators such as Siemens to pursue automation research in the first place depends on the wages that their potential customers face in different countries.¹⁴ To link patents with their owners, we use Orbis Intellectual Property.¹⁵

In our baseline regression, we assume that a firm's innovation in automation is given by the following Poisson specification:

$$PAT_{Aut,i,t} = \exp \left(\begin{aligned} &\beta_{w_L} \ln w_{L,i,t-2} + \beta_{w_H} \ln w_{H,i,t-2} + \beta_X X_{i,t-2} + \beta_{Ka} \ln K_{Aut,i,t-2} \\ &+ \beta_{Ko} \ln K_{other,i,t-2} + \beta_{Sa} \ln SPILL_{Aut,i,t-2} + \beta_{So} \ln SPILL_{other,i,t-2} + \delta_i + \delta_{j,t} \end{aligned} \right) + \epsilon_{i,t}. \quad (3)$$

$PAT_{Aut,i,t}$ denotes the number of automation patents applied for by firm i in year t . $w_{L,i,t-2}$ and $w_{H,i,t-2}$ denote the average low-skill and high-skill wages (more generally labor costs) faced by the customers of firm i at time $t-2$ (we explain below how we proxy for them). Section 3 predicts that $\beta_{w_L} > 0$: an increase in the average low-skill wage faced by the customers of firm i leads firm i to undertake more automation innovations. It also predicts that $\beta_{w_H} < 0$ since high-skill workers are complements to machines. $X_{i,t}$ represents a vector of additional controls (average GDP per capita, GDP gap and labor productivity). Labor productivity can capture technology or human capital shocks in the country where machines may be sold, GDP per capita can capture similar shocks but also demand shocks and the GDP gap business cycles fluctuations and changes in demand.

Ideally, one may want to measure the cost of labor for automatable tasks or occupations (as identified by Webb, 2000, for instance) instead of low-skill and high-skill workers. Unfortunately, in the absence of good international occupational labor costs data, we cannot pursue this approach. Insofar as low-skill and middle-skill workers are those whose tasks have been more intensely automated, our low-skill wage measure can be used as a proxy for the cost of automatable tasks. This proxy will be particularly good

¹⁴If the automation innovation is internal to the firm, then the argument follows if one interprets the innovator's customers as the different downstream production sites of the same firm.

¹⁵For companies in the same business group, R&D decisions could happen at the group level, though treating a group as one agent is often too aggressive (for instance because subsidiaries might be in different sectors). Therefore, for firms within the same business group, we normalize company names by removing non-firm specific words such as country names or legal entity types and then merge firms with the same normalized name. All other firms are treated as separate entities. Therefore, Siemens S.A., Siemens Ltd. or Belgian Siemens S.A. are merged, but Primetals Technologies Germany GmbH which belongs to the same group remains a separate entity in our regressions.

if labor markets are flexible across occupations within education groups or if the labor shocks which move low-skill wages affect low-skill workers similarly across occupations. Otherwise, our use of a noisy measure should result in downward bias.

Following ADHMV, we include controls for knowledge stocks at the firm and country level. $K_{Aut,i,t-2}$ and $K_{other,i,t-2}$ denote the stocks of knowledge in automation and in other technologies of firm i at time $t - 2$. These knowledge stocks are computed using the perpetual inventory method.¹⁶ $SPILL_{Aut,i,t-2}$ and $SPILL_{other,i,t-2}$ similarly denote the stocks of external knowledge (spillovers) in automation and in other technologies which firm i has access to at time $t - 2$ (we explain below how these are constructed). These controls ensure that we are not simply capturing the fact that some firms or countries are on different automation trends. δ_i are firm fixed effects. $\delta_{j,t}$ are industry-year fixed effects (in some specifications we only have year fixed effects). The industry j of a firm is the industry of manufacturing and corresponds to its 2 digit industry in Orbis. Appendix Table A.2 gives the distribution of firms and patents across the main industries in our sample. Finally, $\epsilon_{i,t}$ is an error term. The right-hand side variables are lagged by 2 years in the baseline regressions to reflect the delay between changes in R&D investments and patent applications—Section 5.6 considers alternative timing assumptions.¹⁷ We use the `ppmlhdfe` command from Correia, Guimaraes and Zylkin (2020), which allows to run Poisson regression models with high dimensional fixed effects.

4.2 Macroeconomic data

Our macroeconomic variables come primarily from the 2013 release of the World Input Output Tables, henceforth, WIOD (Timmer et al. 2015). The database contains information on hourly labor costs across groups of educational attainment (low-, middle- and high-skill workers) for the manufacturing sector from 1995 to 2009 for 40 countries including all major markets (US, Japan, all EU countries of 2009, China, India, Brazil, Russia, etc.). We get similar data from the Swiss Federal Statistical Office to

¹⁶We use $\ln(1 + K)$, a depreciation rate of 15% and a dummy for whether the knowledge stock is 0.

¹⁷To control for firm-level fixed effects, our baseline specification uses the Hausman, Hall and Griliches (1984, HHG) method which is the count data equivalent to the within-group estimator. Technically, this method is inconsistent with equation (3) as it requires strict exogeneity and hence prevents the lagged dependent variable from appearing on the right-hand side (which it does through the knowledge stock $K_{Aut,i,t-2}$). Yet, we show in Section 5.6, that our coefficients of interest are not affected by Nickell’s bias by either removing the stock control or by implementing the Blundell, Griffith and Van Reenen (1999) method, which uses the pre-sample average of the dependent variable to proxy for the fixed effect, in line with the patent literature.

add Switzerland, a large source of patents, to our analysis. For our baseline regressions, we focus on labor costs in manufacturing but check that our results are robust to using labor costs in the entire economy. Although our measures cover all labor costs, we refer to those as wages for simplicity. From the same dataset, we obtain measures of labor productivity (as value added divided by hours) and producer price indices (PPI for the whole economy and manufacturing). We obtain exchange rate and GDP data from UNSTAT and compute the GDP gap to control for business cycles.¹⁸ All macroeconomic variables are deflated in the same way. In the baseline regression, we first deflate nominal values by the local PPI for manufacturing (indexed to 1995), and then convert everything into dollars using the average exchange rate for 1995 the starting year of our regressions. Appendix A.5 provides further details.

In the data, low-skill workers are defined as those without a high-school diploma or equivalent and high-skill workers as those with at least a college degree. Middle- and low-skill wages are very highly correlated so one should interpret our low-skill wage variable as reflecting both.¹⁹

The countries with the highest low-skill wages (actually labor costs in manufacturing) in 2009 are Belgium, Sweden and Finland with \$41.9, \$42.2 and \$43.6 respectively (in 1995 dollars) and those with the lowest are India, Mexico and Bulgaria with \$0.28, \$0.61 and \$0.71, respectively. The corresponding number for the US is \$13.7. Table 4 summarizes these values and further shows that the ratio of high-skill to low-skill wages varies considerably across countries, even among those that have relatively similar low-skill wages. The skill-premium in the United States rose from 2.46 to 3.02 during this period while it slightly declined in Belgium from 1.56 to 1.46.

4.3 Computing firm’s market-specific wages and spillovers

Ideally, we would like to measure the wages paid by the (actual and potential) customers of automation innovators. Since we do not observe them, we build a proxy which is a weighted average of country-level wages where the weights reflect the market exposure of innovators. We define the average low-skill wage faced by a firm’s customers $w_{L,i,t}$ as

¹⁸We use a HP filter with a smoothing parameter of 6.25 on $\ln(GDP)$ to get the trend, and the GDP gap is measured as the difference between $\ln(GDP)$ and its trend.

¹⁹For our baseline sample of firms, included in Table 6 below, the correlation between low-skill and middle-skill wages is 0.94 controlling for firm and industry-year fixed effects versus 0.6 for low-skill and high-skill wages. See Appendix Table A.3.

Table 4: Low-skill wages and the skill-premium in manufacturing for selected countries

Country	Low-skill wages (1995\$)		High-skill wages (1995\$)		Skill premium (HSW/LSW)	
	1995	2009	1995	2009	1995	2009
India	0.19	0.28	0.89	1.38	4.79	4.98
Mexico	0.89	0.61	3.46	2.56	3.90	4.21
Bulgaria	1.29	0.71	4.27	1.60	3.32	2.25
United States	11.57	13.67	28.42	41.23	2.46	3.02
Belgium	29.50	41.89	45.98	61.24	1.56	1.46
Sweden	19.92	42.16	34.44	55.92	1.73	1.33
Finland	23.41	43.63	28.10	63.71	1.20	1.46

Note: Wages data, taken from WIOD. The table shows manufacturing low-skill and high-skill wages (technically labor costs) deflated by (manufacturing) PPI and converted to USD using average 1995 exchange rates. Skill-premium is the ratio of high-skill to low-skill wages. The table shows the three countries with the lowest low-skill wages in 2009, the three with the highest and the US.

$$w_{L,i,t} \equiv \sum_c \omega_{i,c} w_{L,c,t}, \quad (4)$$

where $w_{L,c,t}$ is the low-skill wage in country c at time t and $\omega_{i,c}$ is the fixed weight of country c for firm i . We use the same approach to compute average high-skill wages, productivity or GDP per capita. Firms have different exposure to different markets because of trade barriers, heterogeneous tastes of customers, or various historical accidents if exporting involves sunk costs. This is a shift-share measure. The weights are computed pre-sample to ensure that they are weakly exogenous as patent location could be influenced by innovation shocks. Since the weights are fixed, our identification relies on how country-level shocks affect firms differently. In fact, had we observed the wages of the customers of automation innovators, those would have suffered from reverse causality, and we would have used our measure as an instrument. Our regression should therefore be viewed as the reduced form of this instrumental approach. We discuss the recent literature on shift-share regressions in detail in Section 5.5.²⁰

To measure the weights in the absence of sales data, we follow and expand on the methodology of Aghion et al. (2016, ADHMV). We use the firm's pre-sample history of patent filing as a proxy for its market exposure. A patent grants its holder the exclusive right to commercially exploit a technology in a specific country for a limited period of time and inventors must file a patent in each country where they wish to protect their

²⁰As we keep the weights fixed we look at how wage changes in the countries where a firm already sells affect the firm's automation innovation. A different question would have been to analyze how wage changes affect a firm's decision to enter a new market. This is beyond the scope of this paper.

technology. Patenting is costly: a firm needs to hire lawyers and possibly translators as well as pay the filing costs. Further, the publication of a patent can increase vulnerability to imitation and inventors are therefore unlikely to apply for patent protection in a country unless they are relatively certain of the potential market value for the technology (Eaton and Kortum, 1996). Indeed, empirical evidence suggests that inventors do not patent widely and indiscriminately, with the average invention only patented in two countries (see Dechezleprêtre et al., 2011).

We compute for each firm the fraction of its patents in machinery protected in each country c for which we have wage data, $\tilde{\omega}_{i,c}$, during the pre-sample period 1970-1994. We restrict attention to patent families with at least one citation (not counting self-citations) to exclude the lowest quality patents. See Appendix A.6 for details notably on EPO patents. Patenting indicates whether the firm intends to sell in that market. However, a patent in Belgium and one in the U.S. are unlikely to reflect the same market size. At the same time, a larger market attracts more firms so that the market size per firm will generally not grow 1 for 1 with country size. To account for this we weigh each market c by $GDP_{0,c}^{0.35}$, where $GDP_{0,c}$ is the 5 year average GDP of country c at the end of the pre-sample period.²¹ As a result, the weight of country c for firm i is given by:

$$\omega_{i,c} = \frac{\tilde{\omega}_{i,c} GDP_{0,c}^{0.35}}{\sum_{c'} \tilde{\omega}_{i,c'} GDP_{0,c'}^{0.35}}.$$

We use alternative weighting schemes in Section 5.6.

ADHMV verify that a similar method accounts well for the sales distribution of major auto manufacturers. Coelli, Moxnes and Ulltveit-Moe (2020) carry out a more systematic exercise and verify such a method accounts well for aggregate bilateral trade flows and firm exports across 8 country groups in a representative panel of 15,000 firms from 7 European countries (regressing patent weights on sales weights gives a coefficient of 0.89 with a s.e. of 0.008). In supplemental material available on our web page, we also show that our patent weights correlate well with trade flows.

Given that knowledge spillovers have a geographical component, we use the location of firms' innovators to build a measure of the stock of knowledge to which a firm is exposed. We follow ADHMV and compute the stocks of automation patents and of

²¹Eaton, Kortum and Kramarz (2011) estimate the elasticity of French exports to GDP of the destination country to be 1 and the elasticity of the number of French exporters to be 0.65. This gives an elasticity of the average export by firm of 0.35. ADHMV use a power of 1 on GDP instead of 0.35.

Table 5: Descriptive statistics for firms in our baseline regression

Variable	Auto95		Auto90		Weights	Auto95	Auto90
	per year	1997-2011	per year	1997-2011			
Automation patents							
Mean	0.78	11.65	0.92	13.79	Largest country	0.47	0.46
Standard deviation	3.97	52.60	4.71	62.55	Second largest	0.17	0.18
p50	0	2	0	2	US	0.21	0.21
p75	0	6	0	7	Japan	0.17	0.15
p90	1	19	2	22	Germany	0.2	0.21
p95	3	42	4	49	France	0.08	0.09
p99	13	184	15	216	UK	0.09	0.09
Number of firms	3341		4903				

other patents in each country. Then, for each firm, we build a weighted average of country-level knowledge stocks, where the weights correspond to the location of their innovators pre-sample in 1970-1994.²²

4.4 Descriptive statistics

Our basic dataset consists of applicants who have applied to at least one biadic automation patent between 1997 and 2011, who have at least one patent prior to 1995 which can be used to compute weights, and who are not fully domestic (we exclude firms which have only patented in one country pre-sample). For the auto95 measure this corresponds to 3,341 firms, which are responsible for 35,803 or 58% of the total number of biadic auto95 innovations. Table 5 gives some descriptive statistics on the number of automation patents per year and the country weights for the firms in our sample. Over the period 1997-2011, the median firm in the sample filed 2 auto95 and auto90 patent applications. The distribution is very skewed and the 99th percentile firm in the sample has filed 184 auto95 patents. The largest country for a given firm has on average a weight of 0.47 (for auto95). To ensure that our results are not driven solely by the largest country, which we refer to as the “home country” of a firm, we will include in some regressions home country-year fixed effects. The second largest country has on average a weight of 0.17. The three countries with the largest weights on average are the United States, Germany and Japan. Appendix Table A.4 gives the list of the ten biggest automation patenters in our sample.

²²The country stocks are built using the perpetual inventory method with a depreciation rate of 15%. We add dummy variables indicating when the spillover stocks are zero.

5 Global Wages and Induced Automation

We present our main results in three steps: First, our baseline regressions use the full variation of firm low-skill wages to estimate the effect of an increase in low-skill wages on automation innovations. Second, we use country-year fixed effects to isolate the contribution of foreign wages. Third, we contrast the results on automation innovations with those on other types of machinery innovations. The rest of the section discusses identification assumptions, contains robustness checks and additional results including on the minimum wage.

5.1 Baseline results

Table 6 contains our baseline results. The dependent variable is the number of automation (auto95) biadic patents. We use the years 1997-2011 for the dependent variable and, due to the lag structure, 1995-2009 for the independent variables. Recall that skill-dependent wages are measured in the manufacturing sector and deflated by the PPI in that sector.

Table 6: Baseline regressions: effect of wages on automation innovations (auto95)

Dependent variable	Auto95								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.21*** (0.51)	2.83*** (0.73)	1.80** (0.74)	2.46*** (0.75)	2.32*** (0.81)	2.54*** (0.84)	2.90*** (0.78)	2.67*** (0.84)	3.59*** (0.94)
High-skill wage		-0.92 (0.71)	-0.88 (0.67)	-1.56** (0.64)	-1.73** (0.72)	-1.43** (0.71)	-2.22*** (0.72)	-2.61*** (0.79)	-1.52* (0.80)
Stock automation			-0.13** (0.05)	-0.13*** (0.05)	-0.13*** (0.05)	-0.13*** (0.05)	-0.15*** (0.05)	-0.15*** (0.05)	-0.15*** (0.05)
Stock other			0.63*** (0.06)	0.64*** (0.05)	0.64*** (0.05)	0.64*** (0.05)	0.65*** (0.05)	0.65*** (0.05)	0.65*** (0.05)
GDP gap				-3.30 (2.52)	-3.50 (2.59)	-3.00 (2.65)	-4.30* (2.57)	-4.82* (2.66)	-2.69 (2.74)
Labor productivity					0.40 (0.92)			0.91 (0.92)	
GDP per capita						-0.31 (1.14)			-1.94 (1.29)
Spillovers automation							0.60* (0.31)	0.63** (0.31)	0.76** (0.32)
Spillovers other							-0.27 (0.22)	-0.32 (0.22)	-0.41* (0.24)
Fixed effects	F+Y	F+Y	F+Y	F+IY	F+IY	F+IY	F+IY	F+IY	F+IY
Observations	50115	50115	50115	49174	49174	49174	49174	49174	49174
Firms	3341	3341	3341	3329	3329	3329	3329	3329	3329

Note: Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson fixed effects regressions (HHG). All regressions include firm and year or year-industry fixed effects. All regressions with stock variables (resp. spillover variables) include a dummy for no stock (resp. no spillover). Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Column (1) shows the results with only firm and year fixed effects. A higher low-skill manufacturing wage for the customers of an innovating firm predicts more automation innovation. The estimated coefficient is an elasticity so that an increase of 1% in the low-skill wage is associated with 2.2% more automation patents. Column (2) introduces high-skill wages as a control. In all specifications, high-skill wages have a negative (though not always significant) coefficient. Column (3) adds controls for the firm’s stock of knowledge: a higher stock of automation knowledge predicts less automation innovations. Column (4) adds industry-year fixed effects and controls for the GDP gap. Columns (5) and (6) add controls for labor productivity in manufacturing and GDP per capita. None of these macroeconomic controls have consistent significant effects. Columns (7) to (9) repeat columns (4) to (6) but include knowledge spillovers and find that firms which are exposed to more knowledge in automation technologies innovate more in automation. In all specifications, the coefficient on low-skill wages is highly significant with elasticities between 1.8 and 2.9 for columns (1) to (8) and a larger elasticity of 3.6 in column (9). From column (4) onwards, we also control for industry-year fixed effects (where the industry is that of the innovating firm).

In the baseline specification, we cluster at the firm-level to account for auto-correlation in errors. As firms in the same country might be affected by common shocks, we cluster standard errors at the home country (i.e. the country of largest weight) level in Appendix Table A.5. If anything, this tends to reduce the standard error on low-skill wages, a pattern that repeats itself throughout the specifications.²³

5.2 Country-year Fixed Effects and Foreign wages

Country-level shocks which we have not controlled for may impact both wages and innovation, by affecting the cost of innovation or the demand for automation machines through other channels than wages. A tax reform in Germany, for instance, could affect both German low-skill wages and the incentive to innovate. Shocks that mainly affect firms through their home country can be captured through home country-year fixed effects in which case our estimation procedure relies on variation in foreign wages.

²³A potential explanation for the negatively correlated error terms, is that a successful innovation by one firm reduces the innovation of its competitors as the market is already captured. In addition, standard errors may overstate confidence levels if the number of clusters is small or the size distribution of clusters is skewed. To address this, Appendix Table A.5 also includes p-values for low-skill wages using the BDM bootstrap-t approach of Cameron, Gelbach and Miller (2008). All coefficients remain strongly significant.

Table 7: Country-year fixed effects

Dependent variable	Auto95					
	Domestic+Foreign			Foreign		
	(1)	(2)	(3)	(4)	(5)	(6)
Low-skill wage	2.21** (0.99)	2.55** (1.13)	3.56*** (1.24)	4.14*** (1.31)	5.08*** (1.54)	4.14** (1.77)
High-skill wage	-2.89*** (0.94)	-2.16** (1.05)	-1.98* (1.05)	-4.29*** (1.29)	-2.95** (1.46)	-4.29*** (1.39)
GDP gap	4.01 (6.85)	4.94 (6.89)	6.31 (7.16)	-0.72 (4.49)	1.29 (4.84)	-0.73 (5.12)
Stock automation	-0.16*** (0.05)	-0.16*** (0.05)	-0.16*** (0.05)	-0.16*** (0.05)	-0.16*** (0.05)	-0.16*** (0.05)
Stock other	0.66*** (0.06)	0.66*** (0.06)	0.66*** (0.05)	0.65*** (0.05)	0.65*** (0.05)	0.65*** (0.05)
Spillovers automation	1.39*** (0.47)	1.38*** (0.47)	1.37*** (0.47)	1.35*** (0.46)	1.32*** (0.46)	1.35*** (0.47)
Spillovers other	-1.07*** (0.36)	-1.04*** (0.37)	-1.09*** (0.36)	-1.09*** (0.35)	-1.08*** (0.35)	-1.09*** (0.35)
Labor productivity		-1.68 (1.76)			-2.15 (1.58)	
GDP per capita			-3.33* (1.88)			0.00 (2.07)
Fixed effects	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	48773	48773	48773	48773	48773	48773
Firms	3324	3324	3324	3324	3324	3324

Note: Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson fixed effects regressions (HHG). All regressions include firm, industry-year, and country-year fixed effects. Columns (4) to (6) use normalized foreign macroeconomic variables. Normalized foreign low-skill wages are defined as the log of foreign low-skill wages interacted with a measure of the importance of foreign markets in the total wage. This measure is computed at the beginning of the sample period and equals the foreign weight times the foreign low-skill wage divided by total low-skill wages. Normalized foreign high-skill wages, labor productivity and GDP per capita are defined similarly. Normalized foreign GDP gap is the foreign GDP gap interacted with the foreign weight. See text for details. All regressions include dummies for no stock and no spillover. Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Moreover country-year fixed effects are useful to address reverse causality: a technology shock that leads German firms to introduce more automation innovations and therefore lower German wages is unlikely to affect non-German wages since each firm is small for foreign countries. Our identification assumption is then that foreign wages are exogenous to the automation innovation of the firm given our set of controls. We discuss a number of potential confounding factors below.

Columns (1) to (3) of Table 7 reproduce Columns (7) to (9) of Table 6 but add country-year fixed effects, where the country of a firm is still defined as the country with the largest weight (using the headquarters' location to define the home country gives similar results). We still obtain a positive effect of low-skill wages on automation innovations with similar elasticities (between 2.2 and 3.6).

Columns (4) to (6) go further and only consider the foreign component of wages (and of the other macroeconomic variables). To do so, we decompose total low-skill wages

$w_{L,i,t}$ into their home and foreign components as $w_{L,i,t} = \omega_{i,D}w_{L,D,t} + \omega_{i,F}w_{L,F,t}$ where $\omega_{i,D}$ is the home weight, $w_{L,D,t}$ the home wage, $\omega_{i,F} = 1 - \omega_{i,D}$ the foreign weight and $w_{L,F,t}$ the average foreign wage. We use the normalized foreign low-skill wage which is defined as $\frac{\omega_{i,F}w_{L,F,0}}{w_{L,i,0}} \log w_{L,F,t}$. The ratio $\frac{\omega_{i,F}w_{L,F,0}}{w_{L,i,0}}$ captures the fact that more internationally exposed firms are more affected by foreign wages and is computed at the beginning of the sample – though we obtain similar results when we use the average value over the whole sample. This specification ensures that our coefficient can be interpreted as an elasticity on total wages: Since $d \log w_{L,i,t} = \frac{\omega_{i,D}w_{L,D,0}}{w_{L,i,0}} d \log w_{L,D,t} + \frac{\omega_{i,F}w_{L,F,0}}{w_{L,i,0}} d \log w_{L,F,t}$, an increase in the normalized low-skill wage by 0.01 corresponds to an increase in total wages by 1% (recall that we have firm fixed effects). Normalized foreign high-skill wages, GDP per capita and labor productivity are defined similarly (as GDP gap is already an average of logs, we directly interact the foreign variables with $\omega_{i,F}$). Once again we find a positive effect of low-skill wages on automation innovation, with somewhat larger elasticities between 4.1 and 5.1. High-skill wages are the only other macro variable with a consistently significant effect, which is negative between -2 and -4.3 . No other paper using the ADHMV methodology controls for country-year fixed effects and separates foreign wages.

Appendix Table A.6 reproduces the regressions of columns (7) to (9) in Table 6 and of Table 7 but for the auto90 measure of automation. The results are very similar but the coefficients on low-skill wages tend to be of a smaller magnitude, in line with auto95 being a stricter measure of automation. This also helps explain the magnitude of our elasticities in the previous tables: our analysis focuses on innovations with a high automation content (and therefore most likely to respond to an increase in wages) and one should not take our estimates directly to measure the average macro response of the economy to an increase in wages.

Skill-premium. In the previous regressions, the coefficients on low-skill and high-skill wages are of a similar magnitude but opposite signs suggesting that a driver of automation innovations is the skill premium. Table 8 directly regresses automation innovation on the log of the inverse of the skill premium. The coefficient on the inverse skill premium is always of the same magnitude as that on low-skill wages in previous specifications and highly significant. To illustrate the magnitude of our coefficients and the effect of spillovers and stock variables, we run a simulation in Appendix A.7 where we uniformly and permanently decrease the global skill-premium by 10%. This increases the average share of automation innovations in machinery by 4.8 p.p. over the time

Table 8: Skill premium

Dependent variable	Auto95								
	Domestic+Foreign						Foreign		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill / High-skill wages	2.49*** (0.69)	2.64*** (0.69)	2.48*** (0.68)	2.56*** (0.87)	2.40*** (0.86)	2.65*** (0.87)	4.25*** (1.25)	4.11*** (1.22)	4.24*** (1.24)
GDP gap	-4.59* (2.55)	-4.86* (2.57)	-4.57* (2.57)	4.24 (6.77)	4.67 (6.71)	5.02 (6.83)	-0.76 (4.50)	-0.25 (4.56)	-0.57 (4.60)
Labor productivity		0.96 (0.64)			-1.28 (1.09)			-0.43 (0.72)	
GDP per capita			-0.04 (0.72)			-1.71 (1.11)			-0.14 (0.88)
Stocks / Spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	F+IY	F+IY	F+IY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	49174	49174	49174	48773	48773	48773	48773	48773	48773
Firms	3329	3329	3329	3324	3324	3324	3324	3324	3324

Note: Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson fixed-effects regressions (HHG). Columns (1)-(3) include firm and industry-year fixed effects. Columns (4)-(9) include firm, industry-year, and country-year fixed effects. Columns (7)-(9) compute the normalized foreign (log) inverse skill premium as the difference between the normalized (log) foreign low-skill wages and the normalized (log) foreign high-skill wages previously defined. In these columns, GDP gap, GDP per capita and labor productivity also correspond to their normalized foreign values. All regressions include dummies for no stock and no spillovers. Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

period, with 2.7 p.p. coming from the adjustment of stocks and spillovers.

5.3 Non-automation innovations

Is the effect of wages on automation innovations specific to automation or does it affect machinery patents in general? To answer this question, we now look at “placebo” regressions. Specifically, we consider the set of machinery patents and exclude any patent which has a technology category with a prevalence measure above the 60th percentile of the distribution of C/IPC 6-digit codes in the machinery (0.2091). We refer to these as “placebo machinery” innovations. We recompute knowledge stocks and spillover variables for these innovations (“own”) and for all innovations except those (“other”). Table 9 reports the results. Columns (1) to (3) correspond to the baseline regressions with firm and industry-year fixed effects. Low-skill wages only have a positive and significant effect in column (3), but even in that case the coefficient is statistically significantly smaller than with automation (and loses significance with other deflators). Columns (4) to (6) repeat the same regressions but add country-year fixed effects and columns (7) to (9) focus on foreign wages. Neither low-skill wages nor any other macroeconomic control variable has an effect on placebo machinery innovations. The sign of low-skill wages even flip in columns (7) to (9).²⁴ We view this exercise as validating both our empirical

²⁴Conditioning on the 60th percentile is not important and we obtain similar results with machinery innovations excluding auto95 or auto90. Further, replacing low-skill and high-skill wages with the skill

Table 9: Non-automation innovations

Dependent variable	Placebo Machinery								
	Domestic+Foreign					Foreign			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	0.51 (0.60)	0.68 (0.69)	1.73** (0.69)	-0.07 (0.80)	-0.03 (0.91)	0.77 (0.95)	-0.64 (1.19)	-1.22 (1.29)	-0.80 (1.23)
High-skill wage	-0.18 (0.71)	0.12 (0.65)	0.79 (0.75)	-0.27 (0.97)	-0.19 (0.91)	0.20 (1.02)	0.33 (1.18)	-0.39 (1.30)	0.24 (1.33)
GDP gap	-3.39** (1.51)	-3.04* (1.59)	-0.07 (1.90)	-1.15 (3.61)	-1.06 (3.64)	0.31 (3.67)	-2.72 (2.62)	-3.87 (2.78)	-3.02 (2.76)
Stock own	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)	0.03 (0.04)
Stock other	0.55*** (0.04)	0.55*** (0.04)	0.55*** (0.04)	0.56*** (0.04)	0.56*** (0.04)	0.56*** (0.04)	0.56*** (0.04)	0.57*** (0.04)	0.56*** (0.04)
Spillovers own	2.63*** (0.40)	2.66*** (0.41)	2.07*** (0.45)	1.39*** (0.50)	1.39*** (0.50)	1.33*** (0.50)	1.39*** (0.50)	1.35*** (0.50)	1.40*** (0.50)
Spillovers other	-2.31*** (0.46)	-2.32*** (0.46)	-1.79*** (0.49)	-1.32** (0.54)	-1.32** (0.54)	-1.26** (0.54)	-1.33** (0.53)	-1.28** (0.54)	-1.33** (0.53)
Labor productivity		-0.66 (0.72)			-0.17 (1.10)			1.18 (1.14)	
GDP per capita			-3.08*** (0.99)			-1.85 (1.35)			0.27 (1.35)
Fixed effects	F+IY	F+IY	F+IY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	114724	114724	114724	114478	114478	114478	114478	114478	114478
Firms	7696	7696	7696	7693	7693	7693	7693	7693	7693

Note: Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson fixed-effects regressions (HHG). Columns (1)-(3) include firm and industry-year fixed effects, while (4)-(9) include firm, industry-year, and country-year fixed effects. In Columns (7)-(9) the macroeconomic variables are the normalized foreign variables previously defined. Stock and spillover variables are calculated with respect to the dependent variable (placebo machinery). Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

approach and our measure of automation.

5.4 Threats to identification

The previous results establish a correlation between the development of automation technology by a firm and the wages faced by its customers. This relationship is persistent and stable to the inclusion of a number of control variables. Adding country-year fixed effects controls for unobservable shocks to the home country and make it unlikely that reverse causality is the driver. Ideally, we are interested in the effect of an increase in wages on the firm's propensity to introduce automation innovation (in the spirit of our model, we would like to identify $\partial\lambda/\partial w_L$). Of course, wages are an equilibrium outcome, but what matters for identification is that they are exogenous to the inventor. Labor market shocks in manufacturing, such as changes to labor supply or labor costs from regulation, demand for labor in other sectors or demographics, present ideal variation from this perspective. In Section 6, we will therefore focus on a specific labor-market premium in these regressions gives insignificant coefficients. (Results not shown).

shock namely the Hartz reforms in Germany. Consequently, threats to identification arise from other foreign shocks which are correlated with wages and other drivers of automation. We now look at these in detail.

Foreign demand shocks in manufacturing. The biggest threat to identification comes from foreign demand shocks in manufacturing which might drive both wages and the demand for automation equipment. Some aspects of this have already been captured by the controls in Table 7 (GDP gap, GDP per capita and labor productivity). We look at additional controls in Table 10. i) Columns (1) and (5) further control for the share of the manufacturing sector built in the same manner as wages in regressions with country-year fixed effects and either total wages or only foreign wages. The manufacturing share has no consistent effect by itself and does not alter the other coefficients significantly. ii) Conversely, increased offshoring in the foreign country might reduce both wages and the willingness to buy automation technology. We construct a measure of offshoring at the country-level based on the methodology of Timmer et al. (2014): the share of foreign value added in the gross value added in manufacturing. Then, as for other variables, we build a firm-specific value by taking a weighted average. As can be seen from Columns (2) and (6), this does not materially alter the coefficients. iii) In addition, the real interest rate covaries with the business cycle and is a potential important determinant of the cost of purchasing equipment. Columns (3) and (7) show that including the real yield on 10-year government bonds does not alter the coefficients much.²⁵

Labor productivity shocks. We already control for overall labor productivity. An additional concern might come from low-skill specific labor productivity shocks such as $\gamma(i)$ in Section 3, but a positive shock to $\gamma(i)$ would be associated with higher wages and less automation innovation and would correspondingly bias our estimates downwards.

Innovation shocks. A recent period of higher than usual automation innovation might leave both wages and the incentive for further innovation low, creating a spurious positive correlation. To address this, we construct a measure of recent innovation in the same manner as we do for the low-skill wages: for each country we compute the number of automation innovations (from our set of firms or others) applied for in the last three years and then build firm-specific measures. We build a similar control for other innovations. The results in Columns (4) and (8) of Table 10 show that our results carry through (the low-skill wage coefficient in column (4) is just at the margin of significance).

Shocks to the inventing firm. Labor costs also affect the inventing firms through

²⁵We get data for 21 countries (AT AU BE CA CH DE DK ES FI FR GB GR IE IT JP KR LU NL PT SE US) from the IMF and the OECD and deflate nominal yields using the manufacturing PPI.

Table 10: Additional controls

Dependent variable	Auto95							
	Domestic+Foreign				Foreign			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Low-skill wage	2.84** (1.41)	2.97*** (1.15)	2.49** (1.14)	2.01 (1.24)	4.58*** (1.75)	4.97*** (1.50)	5.06*** (1.53)	7.04*** (1.72)
High-skill wage	-1.90 (1.18)	-1.46 (0.99)	-1.98* (1.04)	-1.26 (0.99)	-3.24** (1.48)	-3.02** (1.45)	-2.79* (1.44)	-3.97*** (1.46)
GDP gap	4.30 (6.78)	5.04 (6.87)	5.83 (6.98)	7.02 (6.78)	2.88 (5.32)	1.99 (5.27)	2.07 (4.82)	3.21 (4.93)
Labor productivity	-2.19 (2.21)	-2.69 (1.67)	-1.56 (1.78)	-2.45 (1.85)	-1.99 (1.63)	-1.71 (1.54)	-2.37 (1.58)	-5.14** (2.09)
Manufacturing share	3.69 (9.34)				-6.03 (8.72)			
Offshoring		10.33* (5.51)				-1.93 (4.49)		
Long-term interest rate			0.09 (0.11)				-0.03 (0.06)	
Recent innovation own				-2.77** (1.27)				1.24 (0.92)
Recent innovation other				1.79** (0.77)				-0.34 (0.79)
Stocks / Spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	48773	48773	48467	48773	48773	48773	48356	48773
Firms	3324	3324	3299	3324	3324	3324	3294	3324

Note: Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson fixed-effects regressions (HHG). All columns include firm, industry-year, and country-year fixed effects. “Recent innovation own” denotes the log of a weighted average of automation innovations in the customer’s countries in the last 3 years and “recent innovation other” the same value for all other innovations in these countries. In Columns (4) and (8), wages, GDP gap, labor productivity, and the long-term interest rate are recomputed using weights for the limited set of countries for which interest rates are available (see text). In columns (5)-(8) wages, GDP gap, and labor productivity correspond to their normalized foreign values as previously defined. Normalized foreign recent innovation is defined like the normalized foreign low-skill wages as the interaction between the foreign value and a measure of the importance of foreign markets in the total variable. This measure is computed at the beginning of the sample period and equals the foreign weight times the foreign recent innovation divided by total recent innovation. Normalized foreign manufacturing share, offshoring and long-term interest rates are defined similarly to the normalized foreign GDP gap as the interaction between their foreign component and the foreign weight. Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

its production and R&D costs. Country-year fixed effects alleviate this concern as long as production and R&D are concentrated in the home country. For production costs, if a firm serves a foreign market through offshoring instead of exporting, higher foreign low-skill wages in production would increase the price of machines and therefore bias our coefficient on low-skill wages toward 0. For R&D costs, we can address the potential issue by re-building our firm-specific macro variable using weights based on the location of inventors instead of patent offices. Appendix Table A.7 shows that the coefficient on the baseline low-skill wages remains positive and significant but the coefficient on low-skill wages weighted by inventor weights is small and insignificant. These regressions provide an additional placebo test as they treat firms with the same macroeconomic shocks but weigh them differently.

Placebo. Throughout, our coefficient on low-skill wages should be compared to that from regressions with the placebo machinery innovations, which show persistently little

effect from low-skill wages on innovation. Therefore, if our result on the effect of low-skill wages on automation innovations came from a bias, then that bias would have to be absent for other types of machinery innovations.

5.5 Shift Share

A recent literature addresses the identifying assumptions behind the shift-share set-up in linear regressions. In this respect, it is important for our identification strategy that the weights are pre-determined; that is firms do not choose where to patent based on their expectation of future wages. Appendix Table A.8 demonstrates that country-level growth rates in low- and high-skill wages between 1995 and 2000 have no predictive power on firm weights in 1995. Appendix Table A.9 shows that our results are robust to using weights computed only up to 1989 or to dropping the first 5 years of the regression.²⁶

We interpret our results through the lens of Borusyak, Hull and Jaravel (2018) who show, in the language of our setting, that the random assignment of wage shocks conditional on weights can be sufficient for identification. The inference is valid if either there is a large number of countries (such that the Herfindahl index tends toward 0) affected by independent shocks (controlling for year and firm fixed effects); or the correlation of shocks within a country decays sufficiently rapidly that a large number of country-years is sufficient (see Appendix A2 in their paper).²⁷ They advise practitioners to use appropriate controls to capture omitted variables. We follow this approach by including a large set of controls and country-year fixed effects in our regressions. They recommend applying the standard error correction of Adão, Kolesár and Morales (2019).²⁸

Adão et. al. (2019) show that applications with the shift-share design often lead to an over-rejection of the null. In the language of our application, the problem arises when the residual errors of firms with similar country-distributions are correlated and is not solved by standard clustering. They derive a formula for correcting standard errors in an OLS, which we cannot use directly since we employ a Poisson estimator. Deriving

²⁶Appendix Table A.9 also shows that the results are robust to dropping the earlier years from the weights. It looks at alternatives to premultiplying our patents weights with $GDP_c^{0.35}$: with no multiplication, multiplying by GDP or by total payment to low-skill workers raised to the power of $0.35, (w_L L)^{0.35}$, which may be a better measure of the potential market for technology designed to automate low-skill work. The results are very similar.

²⁷The Herfindahl index is 0.13 and 0.09 when only foreign weights are included. At the country-year level, the corresponding values are 0.009 and 0.006.

²⁸Goldsmith-Pinkham, Sorking and Swift (2020) show that the shift-share setup is valid if the weights are exogenous. In our context, this would require that no unobserved shock—even one uncorrelated with low-skill wages—can affect automation through the weights, which would be a strong assumption.

Table 11: Monte-Carlo simulations

Dependent variable	Auto95								
	Domestic+Foreign						Foreign		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low-skill wage	2.90 [0.105]	2.67* [0.072]	3.59** [0.015]	2.21*** [0.006]	2.55*** [0.003]	3.56** [0.019]	4.14** [0.018]	5.08** [0.023]	4.14*** [0.000]
High-skill wage	-2.22** [0.046]	-2.61 [0.125]	-1.52 [0.156]	-2.89*** [0.005]	-2.16* [0.065]	-1.98** [0.010]	-4.29* [0.062]	-2.95** [0.029]	-4.29 [0.205]
Stocks / Spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
GDP gap	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labor productivity		Yes			Yes			Yes	
GDP per capita			Yes			Yes			Yes
Fixed effects	F+IY	F+IY	F+IY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	49174	49174	49174	48773	48773	48773	48773	48773	48773
Firms	3329	3329	3329	3324	3324	3324	3324	3324	3324

Note: Marginal effects; P-values in brackets. The independent variables are lagged by two periods. Estimation is done by conditional Poisson regressions fixed-effects (HHG). Columns (1)-(3) include firm and industry-year fixed effects. Columns (4)-(9) add country-year fixed effects. Columns (7)-(9) use the normalized foreign macro variables previously defined. All regressions include controls for stocks and spillovers. P-values are computed by sampling with replacement the entire path of macroeconomic variables for each firm with 1000 draws. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

the corresponding correction for the Poisson estimator is beyond the scope of this paper. Instead we implement a Monte Carlo simulation similar to what they do and show that we do not have the same problem of over-rejection.

Specifically, we replicate the regressions of Columns (7) to (9) in Table 6 and of Table 7. For each firm we keep the automation activity, the stocks of innovations, the spillover variables, as well as the distribution of country-weights based on actual data. For each country, we sample with replacement the entire path of macroeconomics variables (wages, labor productivity, GDP per capita and GDP gap) from the existing set with 1000 draws. Table 11 reports the p-values of the coefficients on low-skill wages and high-skill wages based on the simulated distribution of coefficients. The p-values are not markedly different than the ones obtained assuming the standard normal distribution. In particular, the coefficients of interest on low-skill wages are significant at least at the 10% level (except in column 1 with a p-value of 0.105) and at the 2.5% level when we focus on foreign wages. In the language of Adão et al. (2019) the set of controls soaks up most country-specific shocks affecting the outcome variable and, consequently, no shift-share structure is left in the regression residuals.

Finally, Appendix Table A.10 checks that our results are not driven by a single country by sequentially excluding countries in our preferred set-up (with foreign wages and controlling for labor productivity). Excluding a country means that we treat it like the home country when computing normalized foreign wages. We also include the weight of the excluded country times a year dummy as a control. We successively remove the six

largest countries by average weight (US, JP, DE, GB, FR, IT and ES). The coefficient on low-skill wages always remains negative and significant.²⁹

5.6 Additional Robustness Checks and Results

Timing. We look at alternative lags for the dependent variables in Appendix Table A.11 (we keep a lag of 2 for the stock variables, otherwise the dependent variable would be included in the RHS in the lead and contemporaneous cases). The largest coefficient on low-skill wages is obtained for a 2 year lag. It remains relatively stable without country-year fixed effects, while it is more clearly centered around lag 2 with country-year fixed effects.³⁰ Of course, innovators would not be interested in wages 2 years in the past per se, but only inasmuch as they are indicative of future wages. This is our interpretation throughout of our regressions, with the 2 year lag corresponding roughly to the time spent between an effect on R&D and the first results materialized by a patent application. In Appendix Table A.13, we compute predicted future wages at time $t - 2$ based on an AR(1) process with country-specific trends and find similar results.

Innovation types. We look at other definitions or subcategories of automation innovations in Table 12 which reproduces regressions similar to Column (5) of Table 7 with foreign wages and controlling for labor productivity. Column (1) verifies that the results are not driven by the codes that we added to the definition of the machinery technological field listed in footnote 8 (though, we still exclude the weapons categories). Column (2) presents a laxer definition of automation using the 80th percentile of the distribution of the C/IPC 6 digit codes. The effect of low-skill wages is still positive but smaller than for auto90 or auto95. Columns (3) and (4) show that the results are similar for Automat*90 and Automat*80 patents. Automat*90 patents are those which belong to technological categories with a frequency of only the “automat*” group of keywords above the threshold used to define auto90 and automat*80, robot90 or CNC90 are defined analogously. By definition automat*80 patents are all auto80 but 91.5% of

²⁹Goldsmith-Pinkham, Sorking and Swift (2020) advise carrying out a similar exercise by excluding countries with a large Rotemberg weight. Yet, this requires computing our macrovariables as weighted averages of log country-level variables instead of log of weighted averages of country-level variables. We checked that the six countries with the largest Rotemberg weights are the UK, FR, SE, DE, US and BE. Our results are also robust to excluding Belgium and Sweden.

³⁰Appendix Table A.12 carries out placebo regressions where we regress automation innovation on 5, 10 or 15 year leads of wages. We do not find a significant effect of leading low-skill wages (except a negative effect in one column). As expected, given the large number of coefficients a few of the other coefficients are significant but never in a systematic way across specifications.

Table 12: Innovation categories

Dependent variable	AutoX95	Auto80	Automat*90	Automat*80	Robot90	Robot80	CNC90	CNC80
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Foreign:</i>								
Low-skill wage	5.22*** (1.59)	2.62** (1.29)	6.90*** (2.13)	5.53*** (1.94)	5.83* (3.23)	7.09*** (2.49)	-1.71 (4.01)	-1.15 (3.10)
High-skill wage	-1.57 (1.60)	-1.82 (1.33)	-2.05 (1.92)	-1.96 (1.77)	-0.09 (2.91)	-3.27 (2.32)	5.02 (5.40)	0.89 (3.59)
GDP gap	-0.07 (4.50)	1.01 (3.03)	7.76 (5.08)	3.90 (4.64)	4.83 (8.19)	0.21 (6.77)	-1.11 (11.20)	-0.03 (9.34)
Labor productivity	-3.51** (1.69)	-1.09 (1.22)	-5.16*** (1.87)	-4.12** (1.73)	-6.71** (2.73)	-5.10** (2.21)	-3.20 (4.85)	-0.60 (3.15)
Stocks / Spillovers	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY	F+IY+CY
Observations	46980	96695	32738	48950	15927	23060	7609	13417
Firms	3224	6494	2264	3331	1156	1619	582	987

Marginal effects; Standard errors in parentheses. The independent variables are lagged by two periods. Estimation is done by conditional Poisson regressions fixed-effects (HHG). All regressions include firm fixed effects, industry-year and country-year fixed effects. AutoX95 excludes the C/IPC codes which we added when defining the machinery technological field. Auto80 lowers the threshold to define automation innovation to the 80th percentile of the C/IPC 6 digit distribution. Automat*90 and Automat*80 only count words associated with automat. Robot90 and Robot80 only count words associated with robot. CNC90 and CNC80 words associated with CNC. 90 and 80 refer to the threshold used to delimit patents which is the 90th or the 80th percentile of the distribution of automation keywords for 6 digit C/IPC codes. The macroeconomic variables are the normalized foreign variables previously defined. Stocks and spillovers are computed with respect to the dependent variable. Standard errors are clustered at the firm-level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

them are also auto90. Column (5) and (6) shows that our results extend to robot90 and robot80 patents (which are also all auto95). The results differ for CNC patents in columns (7) and (8) perhaps because the sample size is much smaller.

Minimum wage. Given its policy relevance, we also look at the effect of minimum wages using data on 22 countries in regressions where we replace low-skill wages with the minimum wage in Appendix Table A.14.³¹ We find a positive effect of the minimum wage on automation innovations but the coefficients tend to be smaller in the foreign wage regressions and in one specification the coefficient is insignificant. Therefore low-skill wages are a better predictor of automation than the minimum wage. This is not surprising: first, the minimum wage only captures part of the labor costs, second we focus on automation innovations that often happen in manufacturing where low-skill wages tend to be substantially above it, and third we lose nearly half of our countries. An analysis on automation in service industries might show a stronger relationship.

Long-difference. For most of our regressions, we follow the large patent literature and rely on the Poisson estimator, which best handles the count data nature of our dependent variable. In Appendix Table A.15, we conduct a long-difference estimation.

³¹We use data from the OECD. Importantly, not all countries have government-mandated minimum wages and for some countries, we follow the literature and use sectorally bargained minimum wages. See details in Appendix A.5. We do not use the minimum wage as an instrument for low-skill wages because it would be inconsistent: if low-skill wages are endogenous, then high-skill wages are likely endogenous too so that we would need a second instrument.

To allow for zeros in the number of patents, we use the arcsinh transformation and we construct ten 5-year overlapping differences from our 15 years of data. Panel A focuses on firms which patented at least once over the time period considered (now 1995-2013), mirroring what a Poisson regression would do. We find a positive effect of low-skill wages and a negative effect of high-skill wages, although in some specifications the positive effect of low-skill wage is non-significant (in unreported regressions we find that the inverse skill premium always has a positive and significant effect). The diminished significance of low-skill wages reflects the noisy behavior of one-time patenters: Panel B restricts attention to firms which have patented at least twice and recovers the same results as in our Poisson regressions: the change in low-skill wages has a large and significant positive effect on the change of automation innovations. These results suggest that automation responds to medium-run changes in wages.

Additional robustness checks. Our regressions include the stock of automation innovations and therefore may suffer from Nickell’s bias. Appendix Table A.16 removes this variable or uses the usual method of Blundell, Griffith and van Reenen (1999) method, which proxies for the fixed effect with the firm’s pre-sample average of the dependent variable. We obtain very similar results.

Appendix Table A.17 investigates whether our results are robust when focusing on patents of higher quality and weighs patents by citations. We add to each patent the number of citations received within 5 years normalized by technological field and year of application. The results are weaker with total wages and country-year fixed effects but are very similar to the case without weighing patents in our preferred specification with foreign wages and country-year fixed effects.

Appendix Table A.18 shows that our results (using foreign wages and country-year fixed effects) are robust to using different deflators, converting in USD every year or replacing manufacturing wages by total wages. Firms of different sizes may be on different trends in automation innovation. In Appendix Table A.19, we group firms into four bins according to their number of automation patents in 1995, allow for bin-year fixed effects and find similar results.

6 Event study: the Hartz reforms in Germany

To complement our previous analysis on global wages, we now focus on one specific exogenous labor-market shock: the Hartz reforms. The Hartz reforms were a series of

labor-market reforms in Germany designed from 2002 onward and implemented between January 1st 2003 and January 1st 2005. These reforms were the major macroeconomic shock in Germany at the time. They aimed at reducing unemployment and increasing labor-market flexibility by reforming employment agencies to provide better job-search assistance, deregulating temporary work, offering wage subsidies for hard-to-place workers, reducing or removing social contributions on low-paid jobs and reducing long-term unemployment benefits. The reforms have been widely credited with playing a major role in the remarkable performance of the German labor market since, in particular, for increasing labor supply and improving matching efficiency (see e.g. Krause and Uhlig, 2012). In line with the framework of Section 3, such reforms are predicted to reduce the incentive to automate low-skill labor by reducing labor costs (directly through social contribution and indirectly through an increase in labor supply) but also by allowing for more flexible contracts and reducing the expected cost of vacancies.

We use an analogous approach as before to measure innovation and firm's exposure to international markets, but we exclude German firms as they are likely to have been affected by the Hartz reforms through other channels than the labor costs faced by their customers. We run the following regression, over the years 1997–2014:

$$PAT_{Aut,i,t} = \exp(\beta_{DE} \cdot \delta_t \omega_{i,DE} + \beta_{Ka} \ln K_{Aut,i,t-2} + \beta_{Ko} \ln K_{Other,i,t-2} + \delta_i + \delta_{c,t}) + \epsilon_{i,t}.$$

We keep a 2-year lag to the innovation stocks. As before $PAT_{Aut,i,t}$ is a count of automation patents, $K_{Aut,i,t-2}$ and $K_{Other,i,t-2}$ represent firm knowledge stocks, δ_i a firm fixed effect and $\delta_{c,t}$ a country-year fixed effect. $\omega_{i,DE}$ is the fixed firm weight on Germany and δ_t is a set of year dummies (with 2005 as the excluded year). β_{DE} is the full vector of coefficients of interest which determines by how much more a firm exposed to Germany tends to do more automation patents in a given year relative to 2005.

Figure 5.a reports the results. The value of -2.3 in 2010 means that on average a firm with a German weight of 0.1 (the mean value is 0.106) engaged in a 20% smaller increase in automation innovations between 2005 and 2010 than a firm with no German exposure. From 2000 till 2004, firms highly exposed to Germany increased their propensity to introduce automation innovations, such a pre-trend is not surprising in this context since firms exposed to Germany are by definition exposed to different shocks than others. This trend reversed between 2006 and 2009 and resumed from 2010. This is consistent with the Hartz reform increasing labor supply from 2003 onwards and therefore decreasing the incentive to introduce automation innovations from 2005. 2008 marks the beginning




Figure 5: Effect of German exposure on automation innovations. Panel (a) reports coefficients on the interaction between the German weight and a set of year fixed effects in a Poisson regression of auto95 innovations controlling for a full set of fixed effects and firm innovation stocks with 2153 firms. Panel (b) reports coefficients on the triple interaction between the German weight, a dummy for auto95 innovations and a set of year fixed effects in a Poisson regression of auto95 and other machinery innovations controlling for a full set of fixed effects, firm innovation stocks and the interaction between the German weight and a set of year fixed effects with 6452 firms.

of the Great Recession which had a lower impact on German labor markets than in other countries potentially increasing the relative incentive to undertake automation innovations (with an effect 2 years later).

The previous figure clearly shows that the behavior of firms highly exposed to Germany differs over time from that of other firms. To show that the trends above are specific to automation innovations, we run the following regression:

$$PAT_{k,i,t} = \exp \left(\begin{aligned} &\beta_{DE} \cdot \delta_t \omega_{i,DE} + \beta_{DE}^{aut} \cdot \delta_t \omega_{i,DE} 1_{k=aut} \\ &+ \beta_{Ka} \cdot \delta_k \ln K_{Aut,i,t-2} + \beta_{Ko} \cdot \delta_k \ln K_{Other,i,t-2} + \delta_{k,i} + \delta_{k,c,t} \end{aligned} \right) + \epsilon_{k,i,t}. \quad (5)$$

k denotes the type of an innovation which is either auto95 or another machinery innovation, $\delta_{k,i}$ represents a full set of innovation type firm fixed effects, $\delta_{k,c,t}$ innovation type country year fixed effects and $1_{k=aut}$ is a dummy for an auto95 innovation. Standard errors are clustered at the firm level. β_{DE}^{aut} is the vector of coefficients of interests. For each year, they measure how much exposure to Germany increases the relative propensity to introduce automation innovations instead of other forms of machinery innovations compared to 2005. Figure 5.b reports the results: the pattern is, if anything, more

Table 13: Innovation and exposure to Germany

Dependent variables	Auto 95 and other + low auto				Auto95 and low auto	Auto95, low auto and other mach.
	(1)	(2)	(3)	(4)	(5)	(6)
time trend*dummy auto95*German exposure	0.6035*** (0.2294)	0.6037*** (0.2089)	0.7453** (0.3703)	0.0935*** (0.0345)	0.6331*** (0.2266)	0.6034*** (0.2137)
time trend*dummy auto95*post_2003*German exposure	-1.2012*** (0.3965)	-1.2039*** (0.3776)	-1.2942*** (0.4775)	-0.1791** (0.0727)	-1.2580*** (0.4075)	-1.1875*** (0.3840)
dummy auto95*post_2003*German exposure			-0.6836 (1.0374)			
time trend*dummy low auto*German exposure						-0.0260 (0.1236)
time trend*dummy low auto*post_2003*German exposure						0.0647 (0.1768)
year dummy*German exposure	Y	Y	Y	Y	Y	Y
firm innovation stocks * innovation types	N	Y	Y	Y	Y	Y
firm *innovation types fixed effects	Y	Y	Y	Y	Y	Y
country * year * innovation types fixed effects	Y	Y	Y	Y	Y	Y
Observations	77456	77456	77456	77456	62173	107284
Firms	5427	5427	5427	5427	4350	5427

Note: Marginal effects; Standard errors in parentheses. Estimation is by conditional Poisson regressions fixed-effects (HHG). All regressions control for year dummies times the measure of German exposure, innovation stocks (and dummies for no stocks) times the innovation types, firm innovation types fixed effects and country year innovation types fixed effects. The innovation stocks are lagged by two periods. Innovation types are auto95 and all other machinery innovations (low auto and other machinery together) in columns (1) to (4), auto95 and low auto in column (5), and auto95, low auto and other machinery in column (6). German exposure is measured by the German weights in all regressions except for column (4) where it is replaced by a dummy signaling that the firm is in the top quartile of Germany exposed firms. Standard errors are clustered at the firm-level. * p < 0.1; ** p < 0.05; *** p < 0.01

pronounced than in Figure 5.a.

To formally test that the Hartz reform created a trend break in the relative propensity of firms highly exposed to Germany to introduce automation innovation relative to other machinery innovation, we replace the full set of year fixed-effects δ_t in $\beta_{DE}^{aut} \cdot \delta_t \omega_{i,DE} 1_{k=aut}$ in equation (5) with a time trend $t - 2005$ and a time trend interacted with a post 2005 dummy $(t - 2005)_{t > 2005}$. We focus on the years 2000-2010 to have a panel centered on 2005 and avoid the Great Recession. This exercise is akin to a triple diff except that our treatment is distributed continuously (depending on the exposure of each firm to Germany). Table 13 reports the result. Column (2) corresponds exactly to (5): there is a significant time trend in the effect of German exposure on the relative propensity to carry automation innovation between 2000 and 2005, but this trend sharply reverses in the following five years. Column (1) omits the controls for the stock variables. Columns (3) tests whether there is also a shift in level and does not find one. Column (4) replaces the German weight by a dummy indicating that the firm is in the top quartile of German exposure. Column (5) uses the low-automation innovations of section 5.3 instead of all other machinery innovations. Finally, column (6) considers three types of innovations by separating non-auto95 machinery innovations into the low-automation innovations of the previous columns and the rest. In all cases, the trend break remains with a consistent

magnitude (since the 75th percentile of German weight is 0.16). Overall, this exercise shows that, in line with our theory, the Hartz reforms reduced automation innovation of foreign firms highly exposed to Germany, both in absolute terms and relative to other types of machinery innovation.

7 Conclusion

In this paper, we identify automation patents and thereby provide a new measure of automation. Our measure is available at a highly disaggregated level and covers a broad range of technologies. Further, it predicts a decline in routine tasks across US sectors. We then use our classification to analyze the effect of labor market conditions on automation innovations in machinery. We first use global data and find that automation innovations are very responsive to changes in low-skill wages with elasticities between 2 and 5. We proceed to show that the German Hartz reforms led to a relative decrease in automation innovations by foreign firms with a high exposure to Germany. Though using different variations in the data, both exercises emphasize that automation innovations are much more responsive to changes in labor market conditions than other innovations.

These results suggest that policies which increase labor costs for low-skill workers will lead to an increase in innovations which save on them. Therefore, with endogenous technological change, such policies are likely to be less costly for the economy in terms of overall welfare, but also to introduce additional negative effects for low-skill workers. Our paper provides a building block toward estimating by how much a policy-induced increase in low-skill wages would be undone in a couple of years through innovation.

Future research could also adapt our classification method to automation patents beyond machinery. This would allow for an analysis of automation in the service industry or automation of high-skill tasks through Artificial Intelligence.

References

- Acemoglu, D. and Linn, J. (2004). Market size in innovation: Theory and evidence from the pharmaceutical industry. *Quarterly Journal of Economics*, 119(3):1049–1090.
- Acemoglu, D. and Restrepo, P. (2018a). The race between machine and man: Implications of technology for growth, factor shares and employment. *American Economic Review*, 108(6):1488–1542.

- Acemoglu, D. and Restrepo, P. (2018b). Demographics and automation. NBER Working Paper 24421.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.
- Aghion, P., Bénabou, R., Martin, R., and Roulet, A. (2020). Environmental preferences and technological choices: Is market competition clean or dirty? NBER wp 26921.
- Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 214(1):1–51.
- Andersson, D., Karadja, M., and Prawitz, E. (2020). Mass migration and technological change.
- Aschhoff, B., Crass, D., Cremers, K., Grimpe, C., Rammer, C., Brandes, F., Diaz-Lopez, F., Klein Woolthuis, R., Mayer, M., and Montalvo, C. (2010). European competitiveness in key enabling technologies. Technical report, Centre for European Economic Research (ZEW) and TNO.
- Autor, D. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market. *American Economic Review*, 103(5):1553–1597.
- Autor, D., Dorn, D., Hanson, G., Pisano, G., and Shu, P. (2020). Foreign competition and domestic innovation: Evidence from u.s. patents. *American Economic Review: Insights*, 2(3):357–374.
- Autor, D., Levy, F., and Murnane, R. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Bena, J. and Simintzi, E. (2019). Machines could not compete with Chinese labor: Evidence from U.S. firms’ innovation. mimeo.
- Bessen, J., Goos, M., Salomons, A., and van der Berge, W. (2019). Automatic reaction - what happens to workers at firms that automate ?
- Bessen, J. and Hunt, R. (2007). An empirical look at software patents. *Journal of Economics and Management Strategy*, 16(1):157–189.
- Blundell, R., Griffith, R., and Van Reenen, J. (1999). Market share, market value and innovation in a panel of British manufacturing firms. *The Review of Economic Studies*, 66(3):529–554.

- Borusyak, K., Hull, P., and Jaravel, X. (2018). Quasi-experimental shift-share research designs.
- Cameron, A. C., Gelbach, J., and Miller, D. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.
- Clemens, M., Lewis, E., and Postel, H. (2018). Immigration restrictions as active labor market policy: Evidence from the Mexican bracero exclusion. *American Economic Review*, 108(6):1468–87.
- Coelli, F., Moxnes, A., and Ulltveit-Moe, K. (2020). Better, faster, stronger: Global innovation and trade liberalization. *The Review of Economics and Statistics*.
- Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.
- Danzer, A., Feuerbaum, C., and Gaessler, F. (2020). Labor supply and automation innovation. BGPE DIsussion Paper No 198.
- de Rassenfosse, G., Dernis, H., Guellec, D., Picci, L., and van Pottelsberghe de la Potterie, B. (2013). The worldwide count of priority patents: A new indicator of inventive activity. *Research Policy*, 42(3):720–737.
- Dechezleprêtre, A., Glachant, M., Hascic, I., Johnstone, N., and Ménière, Y. (2011). Invention and transfer of climate change-mitigation technologies: a global analysis. *Review of Environmental Economics and Policy*, 5:109–130.
- Dechezleprêtre, A., Ménière, Y., and Mohnen, M. (2017). International patent families: from application strategies to statistical indicators. *Scientometrics*, 111(2):793–828.
- Doms, M., Dunne, T., and Troske, K. (1997). Workers, Wages, and Technology. *Quarterly Journal of Economics*, 112(1):253–290.
- Eaton, J. and Kortum, S. (1996). Trade in ideas: patenting and productivity in the OECD. *Journal of International Economics*, 40:251–278.
- Eaton, J., Kortum, S., and Kramarz, F. (2011). An anatomy of international trade: evidence from French firms. *Econometrica*, 79(5):1453–1498.
- Gaggl, P. and Wright, G. (2017). A short-run view of what computers do: Evidence from a UK tax incentive. *American Economic Journal: Applied Economics*, 9(3):262–94.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *The American Economic Review*, 110(8):2586–2624.
- Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

- Hausman, J., Hall, B., and Griliches, Z. (1984). Econometric models for count data with an application to the patents-R&D relationship. *Econometrica*, 52(4):909–938.
- Hémous, D. and Olsen, M. (forthcoming). The rise of the machines: Automation, horizontal innovation and income inequality. *American Economic Journal: Macroeconomics*.
- Hornbeck, R. and Naidu, S. (2014). When the levee breaks: Black migration and economic development in the American South. *American Economic Review*, 104(3):963–990.
- Krause, M. and Uhlig, H. (2012). Transitions in the German labor market: Structure and crisis. *Journal of Monetary Economics*, 59(1):64–79.
- Krusell, P., Ohanian, L., Ríos-Rull, J.-V., and Violante, G. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68(5):1029–1053.
- Lewis, E. (2011). Immigration, skill mix and capital skill complementarity. *Quarterly Journal of Economics*, 126 (2):1029–1069.
- Lordan, G. and Neumark, D. (2018). People versus machines: The impact of minimum wages on automatable jobs. *Labour Economics*, 52(40-53).
- Mann, K. and Puettmann, L. (2018). Benign effects of automation: New evidence from patent texts.
- Noailly, J. and Smeets, R. (2015). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. *Journal of Environmental Economics and Management*, 72:15–37.
- Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.
- Timmer, M., Dietzenbacher, E., Los, B., Stehrer, R., and de Vries, G. J. (2015). An illustrated user guide to the world input-output database: the case of global automotive production. *Review of International Economics*, 23(3).
- Timmer, M., Erumban, A. A., Los, B., Stehrer, R., and de Vries, G. (2014). Slicing up global value chains. *Journal of Economic Perspectives*, 28(2):99–118.
- Webb, M. (2020). The impact of artificial intelligence on labor market.
- Zeira, J. (1998). Workers, Machines, and Economic Growth. *Quarterly Journal of Economics*, 113(4):1091–1117.